



INSTITUTE FOR DEFENSE ANALYSES

**Doctrinal Guidelines for Quantitative
Vulnerability Assessments of
Infrastructure-Related Risks
Volume I**

J. Darrell Morgeson, Project Leader
Peter S. Brooks
Deena S. Disraelly
Jeremy L. Erb
Michael L. Neiman
Whitney C. Picard

December 2011

Approved for public release;
distribution is unlimited.

IDA Document D-4477

Log: H 11-001952

INSTITUTE FOR DEFENSE ANALYSES
4850 Mark Center Drive
Alexandria, Virginia 22311-1882



The Institute for Defense Analyses is a non-profit corporation that operates three federally funded research and development centers to provide objective analyses of national security issues, particularly those requiring scientific and technical expertise, and conduct related research on other national challenges.

About This Publication

This work was conducted by the Institute for Defense Analyses (IDA) under contract DASW01-04-C-0003, Task ER-6-2474.02, "Guidance Risk and Risk Management Capabilities Development," for the Department of Homeland Security. The views, opinions, and findings should not be construed as representing the official position of either the Department of Homeland Security or the sponsoring organization.

Acknowledgments

The authors wish to thank Dr. Arthur Fries, Dr. Yevgeniy Kirpichevsky, and Dr. James S. Thomason for their review of the paper.

Copyright Notice

© 2011 Institute for Defense Analyses
4850 Mark Center Drive, Alexandria, Virginia 22311-1882 • (703) 845-2000.

INSTITUTE FOR DEFENSE ANALYSES

IDA Document D-4477

**Doctrinal Guidelines for Quantitative
Vulnerability Assessments of
Infrastructure-Related Risks
Volume I**

J. Darrell Morgeson, Project Leader
Peter S. Brooks
Deena S. Disraelly
Jeremy L. Erb
Michael L. Neiman
Whitney C. Picard

Executive Summary

Homeland Security Presidential Directive (HSPD) 7 “establishes a national policy for Federal departments and agencies to identify and prioritize [United States] critical infrastructure and key resources (CIKR) and to protect them from terrorist attack.”¹ Recognizing the complex challenges of developing a risk-based approach for homeland security, Congress asked the National Research Council of the National Academies to “review how the Department of Homeland Security (DHS) is building its capabilities in risk analysis to inform decision making.”²

The National Academies study team examined “the capability of DHS risk analysis methods to appropriately represent and analyze risks from across the Department’s spectrum of activities and responsibilities, including both terrorist threats and natural disasters,”³ and concluded that the basic risk framework—Risk being a function of threat, vulnerability and consequences—is not properly operationalized in many situations and should be revised. Specifically, they found that while threat, vulnerability, and consequences are the right elements for risk assessments (i.e., nothing is missing), the formula $\text{Risk} = \text{TxVxC}$ is not always a sound method for calculating risk.

To address the challenge identified in the National Academies study, the Office of Infrastructure Protection within the National Protection Programs Directorate, Department of Homeland Security asked the Institute for Defense Analyses (IDA) to provide doctrinal guidelines for operationalizing a framework for quantifying risk, with a specific focus on quantitatively estimating the vulnerability of assets and systems comprising the Nation’s critical infrastructure. The IDA study team focused on vulnerability for three reasons. First, its definition and how it is applied to critical infrastructure is far less understood than the concepts of threat and consequence. Second, developing a sound approach for quantifying vulnerability will improve the methodologies for quantifying risk for critical infrastructure. Third, clearly defining

¹ The full text of *Homeland Security Presidential Directive 7: Critical Infrastructure Identification, Prioritization, and Protection*, can be found on the Department of Homeland Security’s webpage, http://www.dhs.gov/xabout/laws/gc_1214597989952.shtm.

² National Academies, National Research Council (NRC), *Review of the Department of Homeland Security’s Approach to Risk Analysis* (Washington, DC: The National Academies Press, 2010), 2.

³ Ibid.

vulnerability is key to developing commensurate risk metrics across the 18 CIKR sectors.⁴

Vulnerability, when used in a critical infrastructure homeland security context, is difficult to define because (1) various homeland security decision-makers define and use metrics for vulnerability in different ways to support their decisions; and (2) cascading consequences, interdependencies, and systems issues can lead to computational complexities that make the idea and valuation of isolated vulnerabilities of little or no use to decision-makers. The lack of a precise definition for vulnerability is evident in the 2010 edition of the *DHS Risk Lexicon* where two separate definitions are provided for the term.⁵ Without further explanation, one is left to infer that the committee that compiled the lexicon could not reach agreement on the exact meaning of the word. IDA's goal is to bridge this conceptual divide by providing a definition and method of computing vulnerability that is satisfactory to all. In addition, the definitions and methodologies presented enable the user to compute risk in a way that produces commensurate risk metrics across the 18 CIKR sectors.

For many CIKR sectors, vulnerability is currently defined as the probability of success given an attack— $P(S|A)$ ⁶—for a given scenario. The IDA study team does not propose altering this definition or its use within a given sector; rather, it proposes combining the probability of success with consequences to produce an expected value of loss metric that can be used to compare vulnerabilities across multiple sectors, including sectors that do not use a conceptual layered defense model.

To this end, the following doctrinal guidelines for quantitative vulnerability assessments encapsulate IDA's results:

1. Define scenarios (combinations of attack vectors/targets) to specify the critical information necessary to estimate scenario risk and its key parameters—consequences, vulnerability, and threat.
2. Ensure that the needs and limitations of the supported decision-making environment are understood and addressed.

⁴ The 18 critical infrastructure and key resources (CIKR) sectors are: Agriculture and Food; Defense Industrial Base; Energy; Healthcare and Public Health; National Monuments and Icons; Banking and Finance; Water; Chemical; Commercial Facilities; Critical Manufacturing; Dams; Emergency Services; Nuclear Reactors, Materials, and Waste; Information Technology; Communications; Postal and Shipping; Transportation Systems; and Government Facilities).

⁵ Risk Steering Committee, *DHS Risk Lexicon*, 2010 ed. (Washington, DC: Department of Homeland Security, September 2010), 38–39.

⁶ $P(S|A)$ is the conditional probability of success given the attack A (noted by the vertical bar, “|”); it summarizes the essential notion of vulnerability to a threat.

3. Use ratio-scale⁷ metrics to quantify threat, vulnerability, and consequences.⁸
4. Quantify vulnerability estimates as an expected value of loss for a given scenario. Two formulations are possible:
 - a. As a product of consequences and probability that an attack is successful—expressed as $C \times P(S|A)$, or
 - b. As a direct estimate of the expected value of loss when $P(S|A)$ estimates are too complex or abstract.
5. Manage the scenario space so that the scope, scale, and assumptions can be verified as appropriate to the decisions, or modified, if determined to be indefensible.

The following bullets summarize how to implement each guideline:

Implementation of Guideline 1

- Quantitative (probabilistic) risk assessments are performed in the context of well-defined scenarios.
- Scenarios are used for risk assessments for intentional hazards (attacks), natural hazards, and accidents.
- The boundaries of a scenario—time horizon, geographic boundaries, and infrastructure systems—are determined by the decision context and decision support requirements.

Implementation of Guideline 2

- There are two doctrinally sound alternative quantifications for vulnerability: (1) the joint probability of successfully penetrating all defensive layers arrayed against the incident (attack, natural hazard, or accident); and (2) the expected value of consequences given that the scenario occurs.
- Quantifying vulnerability as the joint probability of successfully breaching all defensive layers protecting a target facilitates identifying opportunities for risk mitigation through various investments that improve the effectiveness (i.e., lower the probability of success) of defenses at different layers. Risk is formed by combining the probability of success with consequences and/or threat.

⁷ Examples of ratio scale metrics include length, time, probability, cost.

⁸ Not all consequences can be assessed quantitatively. At the present time, the most common set of consequences that are assessed quantitatively are: mortality, morbidity, and economic impacts.

- Given that the scenario has occurred, conditional risk is the expected value of loss or consequences. Conditional risk can be computed by multiplying the joint probability of successfully penetrating all defensive layers times the consequences of the scenario given a success. Overall risk for a given scenario is computed by multiplying conditional risk times the probability that the scenario causing conditional risk occurs.
- Conditional risk is also an appropriate vulnerability metric when there are too many layers, or the computation of the joint probability is too complex and intractable, or both.

Implementation of Guideline 3

- Essential risk parameters and resulting risk metrics are quantified using ratio scales of measurement. Scenarios should be assigned probabilities or frequencies of occurrence depending on the nature of the incident. This is the threat metric.
- Consequences are commonly measured in four ways: human, economic, mission, and psychological. It may also include other factors such as impact on the environment. Typically, the risk metrics are expressed in terms of dollars and lives lost.
- Ordinal scales can be used to qualitatively assess risk. Subject to important limitations, this may be useful for many decision support activities. It is defensible to compare risk interval values to one another, but their lack of precision precludes mathematical combinations. Thus, using combined results from multiple ordinal or ordinal scales integrated with interval scales to compare complex competing risks where mathematical integration of consequences is required is unsupportable in most cases.

Implementation of Guideline 4

- As discussed for Guideline 2, there are two doctrinally sound representations for vulnerability: (1) the joint probability of successfully penetrating all defensive layers arrayed against the attacker; and (2) the expected value of loss given that the scenario occurs, i.e., conditional risk.
- Both representations for vulnerability— $P(S|A)$ or Conditional Risk—are useful in different decision-making contexts and are used by decision-makers to quantify CIKR vulnerability. If the formulation $P(S|A)$ is used, then it should be multiplied by the consequences of the scenario to yield a metric useable in comparative or combined analyses.

- For complex systems or events, conditional risk may be more easily computed using sophisticated models and simulations, or through the incorporation of real data when it exists, or both.

Implementation of Guideline 5

- Producing vulnerability assessments useful for decision-making requires explicitly managing the scenario space; i.e., identification and selection of scenarios that cause risk for the specific decision context being supported.
- For each risk assessment, the scenarios should be selected or constructed to satisfy the following: (1) the scenarios should span the space of possibilities in a manner reasonable and suitable for the decision-making context; and (2) the scenarios should be mutually exclusive.
- For each risk assessment, the scenarios are assigned probabilities or frequencies. For intentional or accidental hazards, these are the probabilities of occurrence given the boundaries of the scenarios being considered. For natural hazards, these are the frequencies of occurrence.

Contents

1.	Overview	1
A.	Background	1
B.	Objective	3
C.	Doctrine for Quantitative Vulnerability Assessments.....	4
D.	Organization of This Document	5
2.	Quantitative Risk and Vulnerability Assessments of Critical Infrastructure and Key Resources (CIKR)	7
A.	Decision Contexts for Department of Homeland Security (DHS) Risk Management	7
B.	CIKR Vulnerability Assessments Require an Integrated Approach	13
C.	Objectives for Doctrinal Guidelines Addressing Quantitative Vulnerability Estimates for CIKR	18
3.	Doctrinal Guidelines for Quantitative Vulnerability Estimates of CIKR	19
A.	Purpose	19
B.	Doctrinal Guidelines	19
4.	Assessment of Doctrinal Guidelines	33
A.	Cross-Walk with Key Recommendations from National Academies Study.....	33
B.	Conclusion.....	34

Appendices

A.	Estimating Vulnerability for Layered Defenses	A-1
B.	Estimating Vulnerability for Biological Attacks on U.S. Populations	B-1
C.	Estimating Vulnerability in the Information Technology Supply Chain (FOUO)—Published Separately	C-1
D.	Ordinal Scales and Risk Assessment.....	D-1
E.	Illustrations	E-1
F.	References	F-1
G.	Abbreviations	G-1

1. Overview

A. Background

Homeland Security Presidential Directive (HSPD) 7 “establishes a national policy for Federal departments and agencies to identify and prioritize U.S. critical infrastructure and key resources (CIKR) and to protect them from terrorist attack.”¹

The *National Infrastructure Protection Plan (NIPP)*

provides the overarching approach for integrating the Nation’s many CIKR protection initiatives into a single national effort. It sets forth a comprehensive risk management framework and clearly defined roles and responsibilities for the Department of Homeland Security; Federal Sector-Specific Agencies; and other Federal, State, regional, local, tribal, territorial, and private sector partners implementing the NIPP.²

There are 18 CIKR sectors.³ Each of these sectors relies on a government-industry partnership to coordinate homeland security regulations, assessments, and investments. The Department of Homeland Security (DHS) develops guidance and tools to assist each sector in assessing its vulnerabilities and risks.

Recognizing the complex challenges in developing a risk-based approach for homeland security, Congress asked the National Research Council (NRC) of the National Academies to “review how DHS is building its capabilities in risk analysis to inform decision making.”⁴

¹ The full text of *Homeland Security Presidential Directive 7: Critical Infrastructure Identification, Prioritization, and Protection*, can be found on the Department of Homeland Security’s webpage, http://www.dhs.gov/xabout/laws/gc_1214597989952.shtm.

² U.S. Department of Homeland Security, *National Infrastructure Protection Plan–2009* (hereafter referred to as *NIPP 2009*) (Washington, DC, 2009).

³ The 18 critical infrastructure and key resources (CIKR) sectors are: Agriculture and Food; Defense Industrial Base; Energy; Healthcare and Public Health; National Monuments and Icons; Banking and Finance; Water; Chemical; Commercial Facilities; Critical Manufacturing; Dams; Emergency Services; Nuclear Reactors, Materials, and Waste; Information Technology; Communications; Postal and Shipping; Transportation Systems; and Government Facilities).

⁴ National Academies, National Research Council (NRC), *Review of the Department of Homeland Security’s Approach to Risk Analysis* (hereafter *Review*) (Washington, DC: The National Academies Press, 2010), 2.

The tasking for the National Academies states the following:

The study will review how DHS is building its capabilities in risk analysis to inform decision-making. More specifically, the study will address the following tasks:

- a) Evaluate the quality of the current DHS approach to estimating risk and applying those estimates in its many management, planning, and resource allocation (including grant-making) activities, through review of a committee selected sample of models and methods;
- b) Assess the capability of DHS risk analysis methods to appropriately represent and analyze risks from across the Department's spectrum of activities and responsibilities, including both terrorist threats and natural disasters;
- c) Assess the capability of DHS risk analysis methods to support DHS decision-making;
- d) Review the feasibility of creating integrated risk analyses covering the entire DHS program areas, including both terrorist threats and natural disasters, and make recommendations for best practices, including outreach and communications; and
- e) Recommend how DHS can improve its risk analyses and how those analyses can be validated and provide improved decision support.⁵

The study by the National Research Council (referred to hereafter as the National Academies study) recommends five actions DHS should take to improve its ability to conduct risk analyses:

1. Build a strong risk capability and expertise at DHS.
2. Incorporate the Risk = $f(T,V,C)$ framework, fully appreciating the complexity involved with each term in the case of terrorism.
3. Develop a strong social science capability and incorporate the results fully in risk analyses and risk management practices.
4. Build a strong risk culture at DHS.
5. Adopt strong scientific practices and procedures, such as careful documentation, transparency, and independent outside peer review.⁶

⁵ Ibid.

The letters T, V, and C, above, refer to the principal components of risk assessments: threat, vulnerability, and consequences, respectively. The threat component concerns the types, instances, and likelihoods of threats, including both natural hazards and acts of terrorism. The vulnerability component generally refers to how susceptible an infrastructure asset or system is to a specified type of threat. The consequences component refers to qualitative and quantitative measures of the resulting effect.

Because T, V, and C may be functionally interdependent, the National Academies study recommends a more carefully developed analytic formulation for combining these components to quantify risk, as compared with the often used multiplicative formulation: $\text{Risk} = T \times V \times C$. The National Academies study concludes that:

1. The basic risk framework of $\text{Risk} = f(T, V, C)$ used by DHS is sound and in accord with accepted practice in the risk analysis field.
2. DHS' operationalization of that framework—its assessment of individual components of risk and the integration of these components into a measure of risk—is in many cases seriously deficient and is in need of major revision.
3. More attention is urgently needed at DHS to assessing and communicating the assumptions and uncertainties surrounding analyses of risk, particularly those involved with terrorism.⁷

The Office of Infrastructure Protection within the National Protection Programs Directorate, Department of Homeland Security asked IDA to explore some of these challenges.

B. Objective

The objective of this document by the Institute for Defense Analyses (IDA) is to provide doctrinal guidelines for operationalizing a framework for quantifying risk, with a specific focus on quantitatively estimating the vulnerability of assets and systems comprising the nation's critical infrastructure. While these doctrinal guidelines address and incorporate the quantification of threat and consequences, this document does not discuss the methodologies and challenges associated with the quantification of threat and consequences.

The IDA study team focused on vulnerability for three reasons. First, its definition and how it is applied to critical infrastructure is far less understood than the concepts of threat and consequence. Second, a sound approach for quantifying vulnerability will improve the methodologies for quantifying risk for critical infrastructure. Third, defining

⁶ Ibid., 88.

⁷ Ibid., 98.

the concept of vulnerability clearly is key to the development of commensurate risk metrics across the 18 CIKR sectors.

The concept of vulnerability, when used in a critical infrastructure homeland security context, is difficult to define precisely. This is so because (1) various homeland security decision-makers define and use metrics for vulnerability in different ways to support their decisions; and (2) cascading consequences, interdependencies, and *systems* issues can lead to computational complexities that make the idea and valuation of isolated vulnerabilities of little or no use to decision-makers.

The IDA study team seeks to define a set of concepts and computational methods for quantifying vulnerability in a way that the resulting risk calculations, which incorporate these concepts, produce commensurable risk metrics regardless of whether the risks are systemic or isolated, or due to natural hazards or adversarial threats. Several examples are given to illustrate these computations and results. In addition, the document discusses some practices for computing vulnerability that are not mathematically sound.

C. Doctrine for Quantitative Vulnerability Assessments

To compare risks across the entire set of critical infrastructures in a commensurate manner, the assessments must conform to certain guidelines. By adhering to these doctrinal guidelines for estimating the risk parameters, the results of individual assessments can be combined or compared in quantitative and defensible ways.

For many CIKR sectors, vulnerability is currently defined to be the probability of success— $P(S|A)$ ⁸—for a given scenario. Examples using this definition are provided in this document (see Appendix A’s discussion on the use of the layered defense model). The IDA study team does not propose to alter this method or its use within a given sector; rather, it proposes combining the probability of success with consequences to produce an expected value of loss vulnerability metric for use in comparisons of vulnerabilities across multiple sectors, including sectors that do not use a conceptual layered defense model.

The doctrinal guidelines are as follows:

1. Define scenarios (combinations of attack vectors/targets) to specify the critical information necessary to estimate scenario risk and its key parameters—consequences, vulnerability, and threat.
2. Ensure that the needs and limitations of the supported decision-making environment are understood and addressed.

⁸ $P(S|A)$ is the conditional probability of success given the attack A (noted by the vertical bar, “|”); it summarizes the essential notion of vulnerability to a threat.

3. Use ratio-scale⁹ metrics to quantify threat, vulnerability, and consequences.¹⁰
4. Quantify vulnerability estimates as an expected value of loss for a given scenario. Two formulations are possible:
 - a. As a product of consequences and probability that an attack is successful—expressed as $CxP(S|A)$, or
 - b. As a direct estimate of the expected value of loss when $P(S|A)$ estimates are too complex or abstract.
5. Manage the scenario space so that the scope, scale, and assumptions can be verified as appropriate to the decisions, or modified, if determined to be indefensible.

D. Organization of This Document

Chapter 2 discusses the challenges associated with producing quantitative vulnerability assessments of critical infrastructure and reviews the National Academies study on DHS' approach to risk analysis for critical infrastructure. Chapter 3 presents the doctrinal guidelines for quantitative vulnerability assessments. Chapter 4 discusses how the doctrinal guidelines satisfy the main recommendations concerning risk analysis methodology found in the National Academies study.

There are seven appendices. Appendix A discusses the layered defense model. Appendix B discusses how the doctrinal guidelines are implemented in a bio-attack scenario. Appendix C describes some of the complexity in the information technology supply chain scenarios. Appendix D summarizes a literature review on the problems with ordinal scaling in the context of risk. This appendix also discusses how ordinal scales pose well-known problems of cognitive bias and misinterpretation and explains through examples how the aggregation of results using ordinal scales in defining risk compounds uncertainty and provides limited information to decision-makers. Appendix E is the list of illustrations contained in the document. Appendix F is a list of references and Appendix G is a list of abbreviations.

⁹ See Appendix D, "Ordinal Scales and Risk Assessment." Ordinal scales are commonly used in homeland security risk analyses. Appendix D discusses how they are often misused and provides the justification for using ratio scales. Examples of ratio scale metrics include length, time, probability, cost.

¹⁰ Not all consequences can be assessed quantitatively. At the present time, the most common set of consequences that are assessed quantitatively are: mortality, morbidity, and economic impacts.

2. Quantitative Risk and Vulnerability Assessments of Critical Infrastructure and Key Resources (CIKR)

A. Decision Contexts for Department of Homeland Security (DHS) Risk Management

Risk analyses in the context of homeland security are complex both conceptually and operationally for two fundamental reasons. The first is the breadth and diversity of decision-making contexts, with varying time horizons within which decisions must be enacted. The second is the various types of CIKR targets at risk, ranging from isolated assets to networked and interconnected systems. As a result, developing consistent and commensurable risk and vulnerability estimating methods that are suitable for a range of threats is quite challenging.

1. Breadth of Decision-Making Contexts

Figure 1 illustrates the breadth of policy, budget, and mission objectives affected by risk informed decision-making in DHS.¹¹ Policy decisions regarding CIKR assets and systems often span 3+ year time horizons and multiple sectors. Risk assessments inform decisions on budgets, policy priorities, mission objectives, anti-terrorism, all hazards preparedness, emergency preparedness, and prevention. They are also informed by other considerations—e.g., overarching federal budget constraints. Though there may be relatively little actual data, quantitative estimates are possible through the use of carefully structured representative or generic cases (e.g., quantitative estimates of threats, targets, vulnerabilities, consequences, and return on investments) that are informed by expert elicitation. For example, decision-making for a time horizon could be informed by the estimated overall casualties in a major city due to bio-terrorism or a hurricane.

Decision-making for the 1–3 year time horizon determines priorities within each sector based on detailed assessments of the threats and potential consequences. Quantitative assessments are needed to substantiate resource allocations, and to contribute to decision-making when multiple types of consequences, both quantitative and qualitative, are considered. Decision-making for the 13 year time horizon would be

¹¹ NRC, *Review*, 29.

informed by estimates of the consequences of specific types of incidents targeting specific parts of the critical infrastructure. For example, a typical query may ask about the short-term disruptions in several CIKR sectors due to a bio-terrorism attack in New York City.

For very near-term decisions, quantitative assessments for CIKR address local security enhancements. A typical query would ask for the optimal allocation of resources to guards, sensors, and remediation for a bio-terrorism attack in the Wall Street area of New York City.

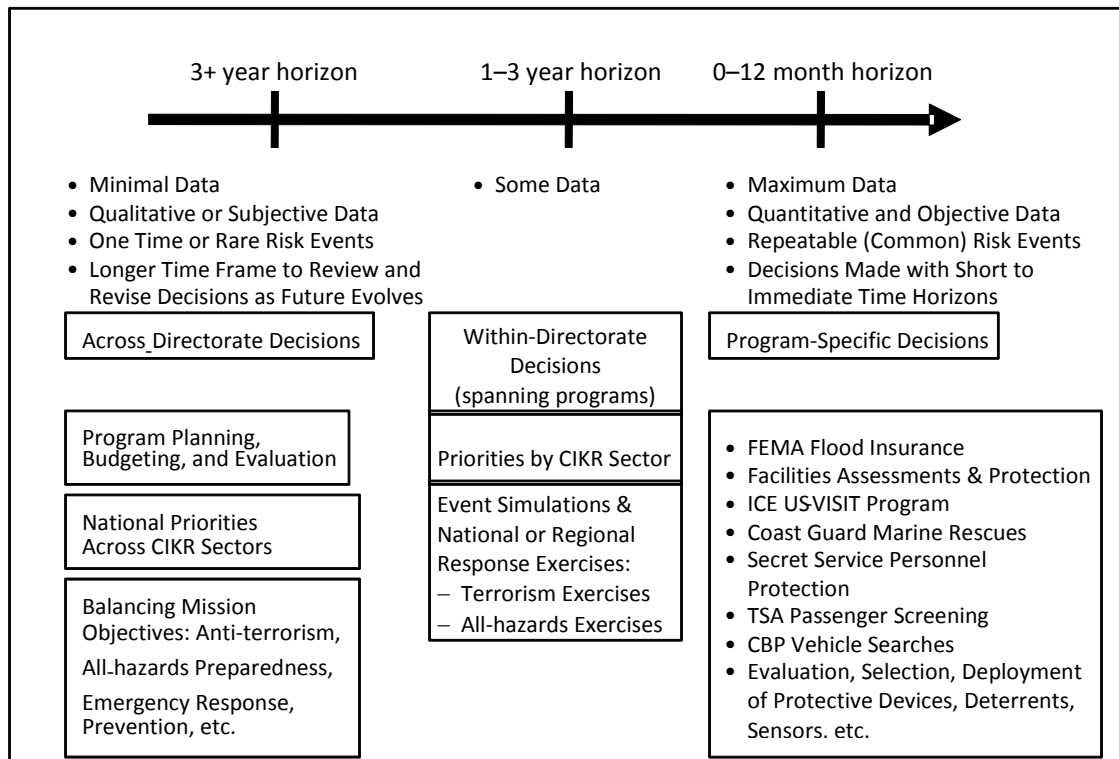


Figure 1. Types of Risk Informed Decisions that DHS Faces (in boxes) Arrayed Roughly According to the Decision-Making Horizon They Inform

In Figure 1, comparative analyses across one or more of the domains may occur at various levels. This requires a common basis of comparison and commensurable metrics. Using scenarios that specify threats and their targets helps structure risk analyses for all the CIKR sectors; and using ratio-scale metrics (such as probabilities, costs, and casualties), ensures that the quantifications of risk can be used to make comparative assessments.

2. CIKR Risk Management

CIKR risk management requires a coordinated effort involving the U.S. Government (USG); the private sector; state, local, and tribal governments; and other organizations. The *NIPP* provides the organizing framework for these partnerships, and is complemented by a series of sector-specific plans (SSPs). These documents guide the development and identification of critical infrastructure assets and systems, the assessment of threats to these targets, and the identification and estimation of vulnerabilities and consequences. The sector partnerships, working with the USG Sector-Specific Agencies, develop initiatives for protecting CIKR resources, and conduct analyses to determine priorities for investments.

The *NIPP* risk management framework consists of six major activities:

- **Set goals and objectives:** Define specific outcomes, conditions, end points, or performance targets that collectively constitute an effective risk management posture.
- **Identify assets, systems, and networks:** Develop an inventory of the assets, systems, and networks, including those located outside the United States, that make up the Nation's CIKR or contribute to the critical functionality therein, and collect information pertinent to risk management that takes into account the fundamental characteristics of each sector.
- **Assess risks:** Evaluate the risk, taking into consideration the potential direct and indirect consequences of a terrorist attack or other hazards (including, as capabilities mature, seasonal changes in the consequences and dependencies and interdependencies associated with each identified asset, system, or network), known vulnerabilities to various potential attack methods or other significant hazards, and general or specific threat information.
- **Prioritize:** Aggregate and compare risk assessment results to: develop an appropriate view of asset, system, and/or network risks and associated mission continuity, where applicable; establish priorities based on risk; and determine protection, resilience, or business continuity initiatives that provide the greatest return on investment for the mitigation of risk.
- **Implement protective programs and resiliency strategies:** Select appropriate actions or programs to reduce or manage the risk identified; identify and provide the resources needed to address priorities.

- **Measure effectiveness:** Use metrics and other evaluation procedures at the appropriate national, State, local, regional, and sector levels to measure progress and assess the effectiveness of the CIKR protection programs.¹²

This framework guides risk management for the 18 CIKR sectors, and is used to prioritize risks and initiatives that “provide the greatest return on investment for the mitigation of risk.”¹³ The relatively independent operation of each CIKR sector contributes to the proliferation of risk assessment models and processes tailored to specific situations. The 18 CIKR sectors and their associated Sector-Specific Agencies are shown in Table 1.¹⁴ The National Academies study lists more than 40 risk assessment models and processes. Another factor contributing to the number of independent approaches is the breadth of assets, threats, and systems and the methodologies to model how each is affected by various threats.

Existing law defines critical infrastructure as “systems and assets” vital to the United States.¹⁵ Asset-based infrastructure sectors are characterized as consisting primarily of physical assets such as commercial structures, factories, and plants. While these sectors may have external resource supply lines and inter- and intra-sector linkages, the primary risks are expressed in terms of the security of individual facilities or a group of closely situated facilities. System-based sectors consist of a collection of assets or facilities that function in concert with one another via a control system; yet in many cases, the assets may be widely dispersed. In system-based sectors, the risks are expressed in terms of critical nodes, system effects, cascading consequences, and consequences that may extend beyond the boundaries of the sector, e.g., affect the general population.

¹² DHS, *NIPP-2009*, 28.

¹³ *Ibid.*

¹⁴ *Ibid.*, 19.

¹⁵ “Critical Infrastructures Protection Act of 2001,” 42 U.S.C. § 5195c(e).

Table 1. Sector-Specific Agencies and Assigned CIKR Sectors

Sector-Specific Agency	Critical Infrastructure and Key Resource Sector
Department of Agriculture Department of Health and Human Services	Agriculture and Food
Department of Defense	Defense Industrial Base
Department of Energy	Energy
Department of Health and Human Services	Healthcare and Public Health
Department of the Interior	National Monuments and Icons
Department of the Treasury	Banking and Finance
Environmental Protection Agency	Water
Department of Homeland Security	
Office of Infrastructure Protection	Chemical Commercial Facilities Critical Manufacturing Dams Emergency Services Nuclear Reactors, Materials, and Waste
Office of Cybersecurity and Communications	Information Technology Communications
Transportation Security Administration	Postal and Shipping
Transportation Security Administration United States Coast Guard	Transportation Systems
Immigration and Customs Enforcement, Federal Protective Service	Government Facilities

Table 2 provides examples of both asset- and systems-based sectors.¹⁶

¹⁶ Department of Homeland Security, *Efforts to Identify Critical Infrastructure Assets and Systems*, DHS OIG-09-86, (Washington, DC: Department of Homeland Security, June 2009).

Table 2. Asset and Systems-based Sectors¹⁷

Asset-based Sectors	Systems-based Sectors
Chemical	Agriculture and Food
Commercial Facilities	Banking and Finance
Dams	Communications
Defense Industrial Base	Energy
Nuclear	Information Technology
<i>The authors denoted the following sectors asset- or systems-based</i>	
Critical Manufacturing	Emergency Services
Government Facilities	Postal and Shipping
National Monuments and Icons	Public Health and Healthcare
	Transportation
	Water

For a fixed facility, such as a chemical manufacturing plant, the processes are engineered, monitored, and controlled by a systems management function. For threats that are physical attacks on the facility, it is common practice to represent the security of the facility as a collection of one or more defensive layers. This is called a layered defense strategy,¹⁸ and gives rise to the concept of vulnerability, defined as the probability that an attack succeeds in breaching all relevant defensive layers given that it occurs. A prime example of a layered defense strategy is perimeter sensors, guards, and gates protecting a building.

For a systems-based sector, the sector's resources are collections of individual assets connected via the transfer of a physical resource, information, and the like. Apart from critical nodes or links, no individual entity may be critical to the functioning of the entire sector. The overall systems management may be dispersed and not under the control of a single authority. While security for such systems addresses the identification and protection of critical nodes and linkages (itself a difficult problem), the greater challenge is to identify and develop countermeasures or responses to systems effects that some threats may cause. An illustrative example is the case of a bio-weapons attack in a city that is not prepared with advanced diagnostics devices or adequate hospital capacity to respond to the event. In this scenario, the different parts of the health care system are not coordinated and can become overwhelmed as the casualties increase. Examples of bio-defense measures include sensors to detect the presence of the bio-agent, early warning medical diagnosis tools, prophylactic measures, and quarantine strategies. The rigorous

¹⁷ Ibid. This report identified 10 sectors as asset- or system-based. The authors provided determinations for the remaining eight sectors.

¹⁸ Appendix A, "Estimating Vulnerability for Layered Defenses."

calculation of the consequences in these cases often involves sophisticated and computationally intensive models that do not formally incorporate a probability of success term as it has been defined for layered defenses. Instead, the focus is on predicting overall human loss, which is often expressed as a probability weighted distribution of losses when taking uncertainties into account (e.g., weather, winds, temperature, populations dispersal, etc.). Appendices A and B illustrate these two examples in more detail.

B. CIKR Vulnerability Assessments Require an Integrated Approach

1. The DHS Approach to Risk Analyses of CIKR

In support of the *NIPP*, DHS and other responsible federal agencies produce threat, vulnerability, and consequence analyses of the nation's critical infrastructure and key resources. These analyses are used by participants in the CIKR sector partnerships to establish priorities for improving the security of the CIKR assets and systems.

The intelligence community, including DHS, conducts threat analyses to identify how various terrorist groups could attack CIKR assets and estimates the likelihood of such attacks. For each CIKR sector, a set of representative scenarios is developed. These are revisited and revised periodically, though the National Academies study highlights the need to do so more formally and methodically, over a broader set of possibilities. This would ameliorate any biases in the process and establish a more complete basis for risk assessments. The National Academies study also recommends ensuring the scenarios and underlying assumptions are more readily understood. This would help decision-makers and users of these analyses understand the probabilities assigned to each scenario and the uncertainties in these estimates caused by either a lack of data or a lack of understanding of threat capabilities.

DHS supports CIKR vulnerability analyses through the work of government and CIKR sector coordinating councils that provide guidance to the asset owners and operators regarding threats and risk methodologies. The owner and operators develop vulnerability assessments; as a result, most vulnerability assessments focus on physical site security. DHS is also pursuing initiatives to address vulnerabilities on a regional or multi-sector level, e.g., the Regional Resiliency Assessment Projects.¹⁹ While there is a growing awareness that vulnerability comprises several facets—physical security, the resilience of a system's response to an incident, and the long-term adaptations of the overall system—current CIKR vulnerability assessment methodologies generally do not take these considerations into account.

¹⁹ NRC, *Review*, 62.

Consequence analyses for CIKR calculate the effects of various threats, measured in economic, health, social disruption, and other terms. The National Infrastructure Simulation and Analysis Center (NISAC), a DHS-funded organization, performs many of these assessments for DHS, modeling, simulating, and analyzing terrorism incidents as well as natural hazards (e.g., hurricanes, earthquakes, floods, and influenza outbreaks). The technical sophistication of the models facilitates the assessments of immediate and long-term effects and also the effects due to cascading influences or networked interdependencies. These analyses can incorporate multiple CIKR sectors, for example, assessing the long-term economic effect in a region when a hurricane damages the electrical grid. An important facet of these analyses is the incorporation of multiple consequence metrics. While these models can address questions of near limitless variation in scope, to be useful, the consequence analyses and thus the risk analyses must be guided first and foremost by the information needs of the decision-makers, including the scope of threats and assets to consider and the time horizon for consequences to be realized.

DHS supports the general view that risk is a function of the above components: threat, vulnerability, and consequences. This is represented by the equation $R = f(T, V, C)$. In practice, there are many independent calculations of the individual risk components T, V, and C, but no uniform method is now used to combine them analytically. Creating a unified method for combining T, V, and C is especially challenging if there is no common basis for what threats are considered, how the CIKR assets and interdependencies are modeled, how vulnerabilities are measured, and what consequences are addressed. For a large majority of CIKR assets, the primary risk analyses concern physical site security assessed with respect to single point threats, e.g., a vehicle-borne improvised explosive device (VBIED) attack against a specific site. For these cases, the multiplicative form— $R = T \times V \times C$ —can be viewed as a specialization of the more general $R = f(T, V, C)$ if T, V, and C are functionally independent, and if the factors T and V are probabilities and C is estimated using ratio scales. For larger or more interconnected CIKR assets, these assumptions regarding the independence of T, V, and C do not hold and more complex models integrating the threat and its effect on the target are necessary.

The use of the multiplicative form, $R = T \times V \times C$, where the three factors are functionally independent, has two additional potential limitations. The factors T and V are, respectively, the probability of the occurrence of the specific threat, and the probability that the threat is successful, if attempted. The meaningful combination of T and V requires that they be expressed using ratio-scale metrics (see Appendix D). Yet it is common, though bad, practice to use non-ratio-scale metrics, e.g., qualitative metrics

(Red, Yellow, Green), or ordinal scales (1, 2, 3, 4, 5) to measure these factors.²⁰ The other potential limitation is the tendency to ignore the uncertainties in the quantitative estimates of T, V, and C. These uncertainties reflect important information about threat capabilities or intentions, or about how the CIKR assets (and the humans involved) will respond during an incident. One example of variations in uncertainty is the intelligence estimate of the likelihood of a specific type of terrorist attack occurring in the short-term versus the long-term. Such uncertainties can affect the determination of overall priorities.

2. Challenges to Quantitative Vulnerability Assessments for CIKR

Vulnerability assessments for critical infrastructure homeland security face several challenges as outlined below:

a. Variety of Decision Contexts for Risk Management Queries

A query from a decision-maker may ask about infrastructure risk in varying contexts, for example (see Figure 1),

- Program planning, budgeting, and evaluation: 3+ year horizon, e.g., to mitigate and deter
- Priorities by CIKR sectors: 1–3 year horizon, e.g., for rapid near to mid-term security upgrades
- Facilities and assets assessments: 0–12 month horizon
- Consequence mitigation: post-event (long enough for cascading consequences to occur).

The decision-maker's use of vulnerability estimates can vary. In some cases, the intent is to determine where to make security investments (i.e., where is the site perimeter most vulnerable?). In others, the intent is to estimate consequences. For example, the query “how vulnerable is New York City to a biological attack” seeks to learn how many casualties would result. The specifics of the decision-maker's query will inform what calculations need to be performed, and how to specify the scope of the analysis.

b. Different Metrics

Various homeland security decision-makers define and use different metrics for vulnerability. While each may be technically defensible, often they are not commensurate; as a result, not usable for comparative evaluations of a set of incidents or a set of investments. A variety of scales are currently in use: ordinal (e.g., Low, Medium, High), interval (e.g., 0 to 10, 10 to 100), and metrics estimated using ratio scales.

²⁰ Appendix D, “Ordinal Scales and Risk Assessment.”

c. Lack of a Consistent Analytic Framework for Quantifying Risk

The quantitative risk metric given that a scenario occurs is called the *conditional risk* associated with that scenario. It is defined as the expected value of loss within a given timeframe should the scenario occur. The expected value of loss from many scenarios can be estimated by incorporating and quantifying the threat metric, which is defined as the probability that the scenario occurs, or $P(A)$. By taking $P(A)$ into account, the total risk for any subset of scenarios within the overall set of scenarios considered in the assessment can be summed to obtain a quantitative estimate of risk overall.²¹ In this discussion, the scenario can be an intentional incident or a natural hazard.²²

For many of the risk calculations for critical infrastructure, the metric of vulnerability, defined as the *probability that an attack is successful given it is attempted*, is computationally defensible and straightforwardly implemented. The typical case involves single point attacks on the physical security of a site. More complex cases can also validly estimate a probability that an attack is successful given it is attempted. In these cases, the conditional risk, defined as the expected value of loss given that the scenario occurs, is calculated as follows:

$$\text{Conditional Risk} = \text{Probability that an attack is successful given that it occurs} \\ \times \text{Consequences of the attack given that it is successful}$$

For other infrastructure elements, the straightforward/literal application of probability of success can become computationally intensive and bewilderingly complex, even when aided by computer simulation. The more useful decision-informing metric for vulnerability in such cases is the conditional risk, calculated as the expected value of loss given that the scenario occurs rather than as the multiplicative product, above. A typical example is a bio-attack in a city—a case in which success may be defined as just the occurrence of the event, on one extreme, or at least 10,000 casualties, for example, on the other extreme. In these cases, the expected value of loss calculation may yield a more informative probability-weighted distribution of casualties with no loss of generality in the overall risk equation.

Even when probabilities of success alone are used to estimate vulnerability, they should be combined with consequences to produce expected value of loss metrics in order to be most informative to decision-makers.

²¹ For this formula to hold, it must be conditional on one, and only one, scenario occurring for the complete set of scenarios being considered in the risk assessment.

²² For natural hazards, probabilities can be ill-suited for events that occur with reasonable frequency during the risk timeframe considered in the analysis. In these cases, it is acceptable to use frequency of events as the estimate of threat (e.g., 20 hurricanes per year).

d. Complex Systems

The complexities of systems (which are comprised of multiple interacting, vulnerable assets), interdependencies, and cascading consequences often overwhelm the capability to produce single-quantity expressions of vulnerability.

For CIKR systems, the possible space of scenario instances can be large due to the large number of scenarios that can be used to attack the system at multiple points. When this is the case, effective management of the scenario space, including aggressive red-teaming efforts, becomes increasingly important. As noted above, consequence calculations, where quantitative metrics are employed, may produce a distribution of results rather than a point estimate. Uncertainties in the probabilities of the scenarios, and uncertainties in the consequence calculations, can be included in the basic calculation of conditional and total risk.²³

For non-quantitative consequence metrics, such as the public's feeling of loss of security, qualitative methods can be used to display the results, provided they are not combined through the misuse of ordinal scales. In cases where a single metric for consequences is desired based on multiple consequence categories, incorporation of utility theory to produce a defensible scale is possible, where there are sound methods for constructing such utility scales.

Another dimension of complexity is the representation of intelligent adversaries. Game-theoretic approaches are one way to combine strategies with the construction of a large set of possible scenario instances. The intertwined, dependent nature of threat, vulnerability, and consequence calculations is a key facet of risk assessments when intelligent adversaries are involved.

e. Imprecise Definition of Success

Yet another challenge to the use of the probability of an attack's success metric (to quantify vulnerability) is defining what success really means. Success for a scenario may actually refer to a range of consequences or a probability-weighted distribution of consequences. For targets that are downstream of the direct effects on a primary target, vulnerability calculations depend on the consequences of the direct attack. This leads to cascading effects for both downstream consequences and vulnerabilities. Both of these analyses are especially important considerations in the calculation of risk to systems. There are cases when the effect of an incident is the denial of service of a particular type, and the system adapts around it. In other cases, the system propagates harm. All of these complexities are considered for a single attack on a specific target. The complexities are

²³ Numerical integration using Monte Carlo computational methods can be used to estimate these distributions.

compounded when multiple, near simultaneous attacks are considered on multiple targets. No prescriptive approach can be applied to decompose these risks. A systems approach allows a consistent method for integrating all these elements of the analyses.

f. Incommensurate Approaches to Risk Assessment

NIPP partners, as participants in a voluntary framework, are not compelled to adopt uniform or consistent and commensurable methodologies that enable cross-sector or cross-threat risk analysis. Risk communication strategies are needed to educate sector partners about the methodologies that facilitate comparative analyses of risks for all CIKR assets. One objective of risk communication strategies is to develop metrics that help explain differences in the public's reaction to terrorism-induced consequences compared to those from natural hazards.

C. Objectives for Doctrinal Guidelines Addressing Quantitative Vulnerability Estimates for CIKR

The National Academies study emphasizes the need to reformulate how T, V, and C are represented in quantitative risk assessments.²⁴

Recommendation: DHS should rapidly change its lingua franca from “Risk = $T \times V \times C$ ” to “Risk = $f(T,V,C)$ ” to emphasize that functional interdependence must be considered in modeling and analysis.

The doctrinal guidelines discussed in the next chapter describe how to operationalize Risk = $f(T,V,C)$ so that the following observations, based on the previously discussed challenges to quantitative vulnerability assessments, are satisfied.

- Risk assessments are valuable insofar as they address a decision-maker's information needs.
- Each risk analysis query has a specified context, i.e., a scope in time, geographic boundaries, threats, assets, and consequence metrics.
- Quantitative risk assessments are calculations of the expected value of loss within a set of scenarios considered.
- The loss calculation considers a scenario instance, defined as a threat instance acting against a target.
- The loss calculation must be expressed in ratio-scale metrics. Other metrics of loss can be considered, but these might not be averaged across a set of cases.

²⁴ NRC, *Review*, 99.

3. Doctrinal Guidelines for Quantitative Vulnerability Estimates of CIKR

A. Purpose

The objective of these doctrinal guidelines is to enable the user to produce sound and commensurable quantitative estimates of the vulnerability of the nation's critical infrastructure systems, assets, and resources. Adhering to these guidelines will allow the user to quantify risks so that they can be effectively compared across the critical infrastructure sectors. In formulating these doctrinal guidelines, the IDA study team focused on quantifying vulnerability, for two reasons: first, to make risk estimates commensurable across sectors that are responsible for a wide diversity of systems, assets, and resources; and second, to gain greater clarity in the fundamental definition of vulnerability among the 18 sectors.

The doctrinal guidelines, concepts, and computational methods in this document will enable the user to calculate vulnerability estimates that are commensurable in side-by-side comparisons. Moreover, they will produce commensurable overall risk metrics, which are critical to support cross-sector comparisons of risk and to enable the objective evaluation of multiple mitigation strategies, and the evaluation of return on investment options. Efforts to quantitatively estimate the vulnerability of assets and systems comprising the nation's critical infrastructure should adhere to the doctrinal guidelines in this document.

In an environment of limited resources and earnest debate about how to manage risks better and more cost-effectively, following these guidelines would contribute to clearer discourse and more efficient use of NIPP partners' resources.

B. Doctrinal Guidelines

The following guidelines provide a more general concept for vulnerability that produces risk metrics for complex scenarios.

1. Define scenarios (combinations of attack vectors/targets) to specify the critical information necessary to estimate scenario risk and its key parameters—consequences, vulnerability, and threat.
2. Ensure that the needs and limitations of the supported decision-making environment are understood and addressed.

3. Use ratio-scale²⁵ metrics to quantify threat, vulnerability, and consequences.²⁶
4. Quantify vulnerability estimates as an expected value of loss for a given scenario. Two formulations are possible:
 - a. As a product of consequences and probability that an attack is successful—expressed as $C \times P(S|A)$, or
 - b. As a direct estimate of the expected value of loss when $P(S|A)$ estimates are too complex or abstract.
5. Manage the scenario space so that the scope, scale, and assumptions can be verified as appropriate to the decisions, or modified, if determined to be indefensible.

The resulting risk metrics are commensurable with risk metrics produced by vulnerability assessments that use the probability of a successful attack.

1. Define Scenarios (Combinations of Attack Vectors/Targets) to Specify the Critical Information Necessary to Estimate Scenario Risk and Its Key Parameters—Consequences, Vulnerability, and Threat.

a. Implementation of Guideline 1

- Quantitative (probabilistic) risk assessments are performed in the context of well-defined scenarios. Scenarios provide the essential context for estimating risk as a function of threat, vulnerability, and consequences.
- Scenarios are used for risk assessments for intentional hazards (attacks), natural hazards, and accidents. These scenarios define the characteristics of postulated incidents.
- Scenarios include the following information: incident, target (or affected infrastructure element), and descriptors of the boundaries of the scenario and analysis in time, geographic boundaries, and infrastructure systems to be included in the risk assessments.
- The terms scenarios, incidents, and targets refer to specific instances. When a collection of similar instances or a representative instance is considered, it is referred to as a type. For example, to facilitate comparison of risks among

²⁵ See Appendix D, “Ordinal Scales and Risk Assessment.” Ordinal scales are commonly used in homeland security risk analyses. Appendix D discusses how they are often misused and provides the justification for using ratio scales. Examples of ratio scale metrics include length, time, probability, cost.

²⁶ Not all consequences can be assessed quantitatively. At the present time, the most common set of consequences that are assessed quantitatively are: mortality, morbidity, and economic impacts.

multiple scenarios, some attack vectors contained in scenarios are generalized/standardized so that they provide a common basis for comparing risk among all targets of interest. When using types, the supporting rationale should be provided to clarify why the characteristics of the type were chosen and how the values for vulnerability and consequences have been estimated.²⁷

- The boundaries of a scenario—time horizon, geographic boundaries, and infrastructure systems to be included in the risk assessments—are determined by the decision context and decision support requirements. The time horizon spans a portion or all of the period when the incident and the consequences are occurring. When comparing risks across a set of scenarios, comparable boundary specifications must be used to ensure commensurate results. Thus, nationally scoped consequences should not be compared to local consequences; a five-year accumulated consequence should not be compared to six months of consequences. The use of stopping rules may facilitate specifying the boundaries, e.g., calculating consequences only up to one year after the incident. These boundaries relate to fair and reasonable comparisons. They preclude comparing a bombing event in a chemical plant and local impact for one year to a bombing event at a shopping mall and the national impact for five years. Nonetheless, they are not meant to limit analysis. If extended boundaries are important to a decision, the case for this should be made, and fair and objective use of the data must be assessed.

b. Discussion

Commonly defined, risk is the potential for an unwanted outcome resulting from an incident, event, or occurrence, as determined by its likelihood and the associated consequences. Quantitatively, risk is estimated as the expected value of loss from one or more scenarios times the likelihood or frequency of those scenarios. A scenario describes an incident (attack, accident, or natural disaster) and what specifically is being attacked or affected. The probability or frequency of a scenario and the associated consequences are calculated with respect to a domain, which defines the extent of the effects calculations. The domain is defined either geographically, or functionally (e.g., cascading, interdependent effects), or both. The domain also includes the time horizon, e.g., immediate, one week, five years.

As an example, consider a set of scenarios comprised of terrorist attacks using IEDs at large commercial buildings on a specified day in New York, and the consequence

²⁷ These estimations can be produced by averaging (a probability-weighted average may be used) the risk for some number of specific instances providing a range of risk estimates, but this may also be performed for scoping purposes with the use of worst-case characteristics.

estimates that focus on the loss of life over the entire metropolitan area and the costs and economic impacts nationally within one year. Another example could be a set of scenarios involving hurricanes that strike the East Coast from June through November of a given year, and the domain could be the national electrical grid and the costs and economic consequences due to the associated large-scale power outages.

To compute a probabilistic risk assessment, a set of scenarios is specified. For each scenario, a probability of occurrence is determined and the expected value of the consequences (EVC) given the scenario occurrence is calculated. The expected value is discussed in section 4, below. The following is a mathematical formula for a probabilistic risk assessment:

Risk(Scenarios S, Domain D) =

$$\sum_{\text{Scenarios } S} P(\text{Scenario } S, \text{Domain } D) * EVC(\text{Scenario } S, \text{Domain } D)$$

2. Ensure that the Needs and Limitations of the Supported Decision-Making Environment Are Understood and Addressed.

a. Implementation of Guideline 2

- There are two doctrinally sound alternative quantifications for vulnerability: (1) the joint probability of successfully penetrating all defensive layers arrayed against the incident (attack, natural hazard, or accident); and (2) the EVC given that the scenario occurs.
- Quantifying vulnerability as the joint probability of successfully breaching all defensive layers protecting a target facilitates identifying opportunities for risk mitigation through various investments that improve the effectiveness (i.e., lower the probability of success) of defenses at different layers.
- For layered defenses, the probability of an adversary successfully breaching the outer layer is not a condition of successfully penetrating any previous layer. For all layers of defense contained within this outer layer, the probabilities are a condition of the sequence of previously penetrated layers. For example, in a two-layer defense—L1 (the outer layer) and L2 (the inner layer)—the joint probability of successfully penetrating both L1 and L2 is $P(L1) \times P(L2|L1)$. That is, the unconditional probability of successfully penetrating L1 is multiplied by the probability of successfully penetrating L2, given that L1 has been successfully penetrated. This general formula can be extended to accommodate multiple layers. However, as the number and complexity of layered defenses increases, this approach can become computationally intractable.

- Vulnerability quantified as probability of success, while useful for determining which layers of defense provide the greatest opportunity for enhancement, is not by itself a meaningful metric to support comparative risk management decision-making. A more meaningful risk metric is formed by combining probability of success with consequences and/or threat to produce the conditional risk or overall risk metric. For example, some scenarios with a high probability of success are not high risk because the consequences may be low or the probability of scenario occurrence may be low (or both), making the product of threat, vulnerability, and consequences not significant when compared with the risk of other scenarios. Hence, in practice, one should avoid making investment or policy decisions based solely on the quantification and comparison of probabilities of success.
- Given that the scenario has occurred, conditional risk is the expected value of loss or consequences. Conditional risk can be computed by multiplying the joint probability of successfully penetrating all defensive layers times the consequences of the scenario, given a success.
- Conditional risk is also an appropriate vulnerability metric when there are too many layers, or the computation of the joint probability is too complex and intractable, or both. Several sectors have developed sophisticated approaches that incorporate expert judgment, historical data from tests and real world events, and computational methods that permit the computation of the expected value of loss directly—that is, by multiplying a vulnerability calculation with the calculation of predicted loss or harm. Some of these methods do not directly compute a probability of success metric; rather they rely on other stronger and more rigorous methods to directly estimate expected consequences given different conditions. Often these methods involve stochastic representations and produce a range or distribution of expected consequences.
- In some cases, both methods for computing vulnerability—probability of success and expected value of loss—are acceptable. The choice of which to use depends on the decision context being supported and the resources (time, funding, etc.) available. Using a biological attack²⁸ as an example, one might consider the targets as the potentially affected population, with varied susceptibility to a pathogen. A defensible analysis might use a sophisticated

²⁸ Biological attacks are infrastructure related risks since they sometimes use infrastructure (food distribution systems, postal and shipping systems, transportation hubs) as a delivery method. If the attack affects a large population, it also affects the function of infrastructure since it relies on qualified people to operate it.

Monte Carlo simulation²⁹ that estimates the outcomes of disease spread by a contagious pathogen, in light of vaccinations, antibiotics administration, the exposed populations' characteristics, early warning, etc. On the other hand, when resources are unavailable or the decision-maker is willing to live with coarser estimates, vulnerability analysis may be set up using an abstract layered defense model (e.g., defensive layers may be described as early warning, protective clothing, or antibiotics). This would be useful for identifying the potential value of mitigations at different layers of defense.

b. Discussion

There are several contexts where vulnerability is used as a decision variable. The first concerns the efficacy of specific defensive measures and an attack that must defeat layers of defense. The decisions concern investments to strengthen individual defensive measures and trade-offs among investing in different types and layers of defense. The second context concerns the susceptibility to overall adverse effects from scenarios acting against complex target sets. The decisions involve complex sets of defensive and mitigation measures. The third context for vulnerability is descriptive versus quantitative, and identifies for a given target or sets of targets what scenarios are likely to cause adverse consequences. In a similar context, decision-makers may need to know which targets are vulnerable to a specific scenario or type of scenario. These are discussed in turn below.

The 2010 edition of the *DHS Risk Lexicon* defines Vulnerability in the following manner:

Vulnerability: Physical feature or operational attribute that renders an entity, asset, system, network, or geographic area open to exploitation or susceptible to a given hazard.

Vulnerability (Degree): Qualitative or quantitative expression of the level to which an entity, asset, system, network, or geographic area is susceptible to harm when it experiences a hazard.³⁰

- A common vulnerability assessment involves identifying how to lower the probability of success for a deliberate attack against a physical asset. A defense comprised of detection, disruption, guards, fences, or other defensive measures that must be penetrated in sequence is called a layered defense. In these cases,

²⁹ Monte Carlo simulations are a class of computational algorithms that rely on repeated random sampling to compute their results.

³⁰ Risk Steering Committee, *DHS Risk Lexicon*, 2010 ed. (Washington, DC: Department of Homeland Security, September 2010), 3839.

vulnerability is defined as the probability that the attack is successful given that it is attempted.

$$\text{Vulnerability} = \text{Probability}(\text{Success}/\text{Attack}) \quad (1)$$

Example: What is the probability that terrorists could successfully penetrate the national layer of defense, enter and operate freely in the country without detection and disruption, position and prepare to attack a facility without local police detecting and disrupting, penetrate the facility's perimeter security, and still cause unacceptable damage or loss?

- For complex targets and attacks, the defenses do not operate as successive layers all engaged before a weapon interacts with a target. The effects, in such cases, can vary over a large and complex domain. For example, the defenses against a biological weapon attack may consist of actions to keep the weapon from being emplaced (layers of physical security), but may also include enhanced resources at hospitals (more doctors, beds, medicine stockpiles) to lessen the transmission of a contagious agent and the severity of the consequences. When the vulnerability query is focused on the overall effects, then

$$\text{Vulnerability} = \text{Expected Value of Consequences (given specified defenses)} \quad (2)$$

The vulnerability query in this case may appear as follows: a decision-maker asks "how vulnerable is the city of New York to an anthrax attack of a specified size, time and place, assuming defenses that lack early warning and limited ability to perform rapid administration of antibiotics post-attack?" The answer "the number of expected lives lost due to an anthrax attack in New York City today would be between 10,000 and 25,000" represents a quantification of the vulnerability using EVC. An investment in early warning and rapid treatment with antibiotics would reduce the expected lives lost to the range of 500 to 1,500, for the same attack scenario. On average, the investment results in a reduction of 16,500 expected lives lost.

- A third context for vulnerability is the identification of infrastructure susceptibilities to exploitation or harm from specified scenarios. The concern "How vulnerable is a given office building to an earthquake of magnitude 6.0?" may seek information about which building strengthening projects have and have not been undertaken.
- A fourth context for vulnerability is the identification of targets that are most vulnerable to a given attack. For example, "Which metro stops are most vulnerable to a 50 lbs. enhanced explosive device?" Notionally, a decision-maker could be interested in those locations and times when the greatest

expected loss of life might occur. Decisions that might be made in such contexts would be to put additional security and surveillance at these locations.

This document provides two methods for quantifying vulnerability. Method 1 quantifies vulnerability as the joint probability of penetrating all defensive layers successfully. It is useful for asset-focused risks where there is decision value in being able to determine what layer of defense against an adversary would provide the most valuable investment, or whether owner/operator investments should be augmented by additional state/local law enforcement in a buffer zone. Method 2 quantifies vulnerability as the expected value of the consequences given a specific scenario occurs. Decision metrics for managing risks due to a biological weapon attack, as one example of a complex system, are best informed by their overall consequences. To address the need for multiple levels of analysis and cost-benefit tradeoffs, this document includes in the vulnerability analysis doctrinal guidelines for specifying and managing the scenario space, and how to work with multiple metrics of risk.

3. Use Ratio-Scale Metrics to Quantify Threat, Vulnerability, and Consequences.

a. Implementation of Guideline 3

- Essential risk parameters and resulting risk metrics are quantified using ratio scales of measurement. Scenarios should be assigned probabilities or frequencies of occurrence based on the nature of the incident. This is the threat metric.
- Consequence is commonly measured in four ways: human, economic, mission, and psychological, but may also include other factors such as impact on the environment. Commonly, the risk metrics are expressed in terms of dollars and lives lost.
- Some consequence metrics can be combined if they can be transformed to a common utility metric. Expected utility theory, used in economics and game theory, may aid the comparison of risk reduction measures that impact different consequence metrics; however, its use is limited because of the difficulty and expense in producing scales for utility that conform to the decision-maker's value system.
- Externally valid metrics (dollars, lives, probabilities) as compared with points, index values, and utilities are more useful within the NIPP, as they reflect objective common values, rather than the unique value system of any given authority.
- Ordinal scales can be used to qualitatively assess risk. This may be useful for many decision support activities, subject to important limitations. It is defensible

to compare risk interval values to one another, but their lack of precision precludes any mathematical combinations. Thus, the use of combined results from multiple ordinal scales or ordinal scales integrated with interval scales to compare complex competing risks where mathematical integration of consequences is required is unsupportable.³¹ In essence, if likelihood³² is estimated using one ordinal scale, and consequences are estimated using another (or several others), there is no meaningful way to produce the combined likelihood and consequence variables. This practice, however, is not uncommon, since it appears to laymen to be intuitive and to follow the risk sciences. However, it violates basic mathematical laws of probability and measurement scales, and is likely to misinform decision-makers.

- Commensurability of all results requires common consequence metrics or utility metrics to combine several consequences into a single parameter. Homeland security decision-makers may want to be able to see quantified expected values of losses according to the type of loss (morbidity, mortality, costs, and economic impacts). In some cases, however, it is desirable to integrate all consequences into a single metric. The ability to compare risks across scenarios also requires the assignment of probabilities or frequencies to each scenario under consideration. The scenarios must have a common specification of the timeframes that risk is being calculated. In addition, comparable specifications for cascading consequences due to interdependencies must be carefully managed to insure comparable risk results.

b. Discussion

The two most defensible ways to estimate vulnerability are as an EVC or as a probability of successfully breaching all defensive layers. The laws and axioms that provide the foundation for the probability theory insure that probability estimates are ratio-scaled. In addition, these estimates of vulnerability together with extended expert elicitation techniques can provide an additional decision metric of potential value to decision-makers—the quantification of uncertainty ranges around point estimates and expected value calculations.

Commensurability of results requires common consequence metrics or utility metrics, and probabilities assigned to each scenario. Commensurability of results also requires the use of ratio scales to represent the probabilities in the risk equation.

³¹ Appendix D, “Ordinal Scales and Risk Assessment.”

³² Probability estimates by definition should not be estimated using ordinal scales.

4. Quantify Vulnerability Estimates as an Expected Value of Loss. Two Formulations Are Possible:

- a) As a product of consequences and probability that an attack is successful, or**
- b) As a direct estimate of the expected value of loss when $P(S|A)$ estimates are too complex or abstract.**

a. Implementation of Guideline 4

- As discussed in Section 2 above, there are two doctrinally sound representations for vulnerability: (1) the joint probability of successfully penetrating all defensive layers arrayed against the attacker; and (2) the expected value of loss given that the scenario occurs.
- For those decisions where there is sufficient confidence in the probability or frequency of the scenario's occurrence, risk—calculated as the product of the probability of a scenario and the conditional risk given a scenario—is the best decision-informing metric.
- Conditional risk can be expressed in several different ways based on the needs of the decision-maker. These include worst-case loss, probability weighted loss, and a bounded range of losses. In some cases, the expected value of loss also may be represented as a continuous probability distribution.
- For complex systems or events, conditional risk may be more easily computed using sophisticated models and simulations, or through the incorporation of real data when it exists, or both. In these situations, conditional risk/expected value of loss given the scenario occurs often becomes the more useful way to quantify vulnerability.
- When conditional risk is used to directly quantify vulnerability, there are three advantages. First, risk analysts can avoid computing the joint probability of successfully penetrating all layered defenses, which may be computationally intractable (e.g., if there are a very large number of layers, as in bio defenses where a defensive layer may be viewed as the immune system of a single individual). Second, rigorously, validated calculations of conditional risk are commonly used for many critical systems and critical events. Significant time and money has been spent to develop powerful computational techniques that directly estimate the expected value of loss given a scenario occurs. Third, decision-makers who use the output of the computational methods are accustomed to using conditional risk directly in their deliberative process and unless they are explicitly trying to compare risk management courses of action

at different layers, would not find value in using a layered defense/probability of success approach.

b. Discussion

To form meaningful comparisons across multiple sectors, the consequence metrics must be of the same type (e.g., economic loss or loss of life). Utility Theory can be used to combine multiple consequence categories into a single utility metric.

The EVC calculation is the probability weighted average—over a set of scenarios—of the estimated consequences associated with each scenario. Consequences can be expressed in several ways, including Worst Case, Bounded Range or a distribution which are described as follows:

- **Worst Case.** This is biased in favor of the attacker by considering scenarios that produce the worst case values of consequences.
- **Bounded Range.** This constructs the worst case as above (favorable to the attacker) and the best case (favorable to the defender) with the expected value falling within the range.
- **Distribution.** This is either continuous or discrete, and is typically produced by expert judgment, historical data, Monte Carlo simulations, or other defensible means.

5. Manage the Scenario Space So that the Scope, Scale, and Assumptions can be Verified as Appropriate to the Decisions, or Modified, If Determined to be Indefensible.

a. Implementation of Guideline 5

- Producing vulnerability assessments useful for decision-making requires explicitly managing the scenario space, i.e., identifying and selecting scenarios that cause risk for the specific decision context being supported.
- For each risk assessment, the scenarios should be selected or constructed to satisfy the following: (1) the scenarios should span the space of possibilities in a manner reasonable and suitable for the decision-making context; and (2) the scenarios should be mutually exclusive. For example, a risk assessment of terrorists attacking New York City using a 500-lb. vehicle-borne bomb should include a set of possible target locations that span the geographic boundaries under consideration. It encourages the selection of potential targets throughout the region, not just those targets that appear attractive.

- For each risk assessment, the scenarios are assigned probabilities or frequencies. For intentional or accidental hazards, these are the probabilities of occurrence given the boundaries of the scenarios being considered. For natural hazards, these are the frequencies of occurrence. By stating the timeframe in which the scenario occurs, both intentional and natural hazards can be compared on a common basis.
- For intentional hazards, the calculation of risk assumes an intelligent adversary that decides what, when, where, and how to attack based on some reasonably assumed knowledge of existing defenses.

b. Discussion

Figure 2 is a top-level architecture for managing the scenario space. There are three main components in this architecture:

1. A representative set of risk management queries.
2. Based on these queries, a robust set of cases (see Figure 2) and the ability to manage the scenario space. This involves selecting scenarios and domains that are tailored to each query. The consequence calculations may be pre-computed or performed on an as-needed basis.
3. The analysis loop “query → case specification → risk calculation” may involve several iterations to reflect the decision-making of an intelligent adversary and to perform analytic excursions to measure different risk management decisions.

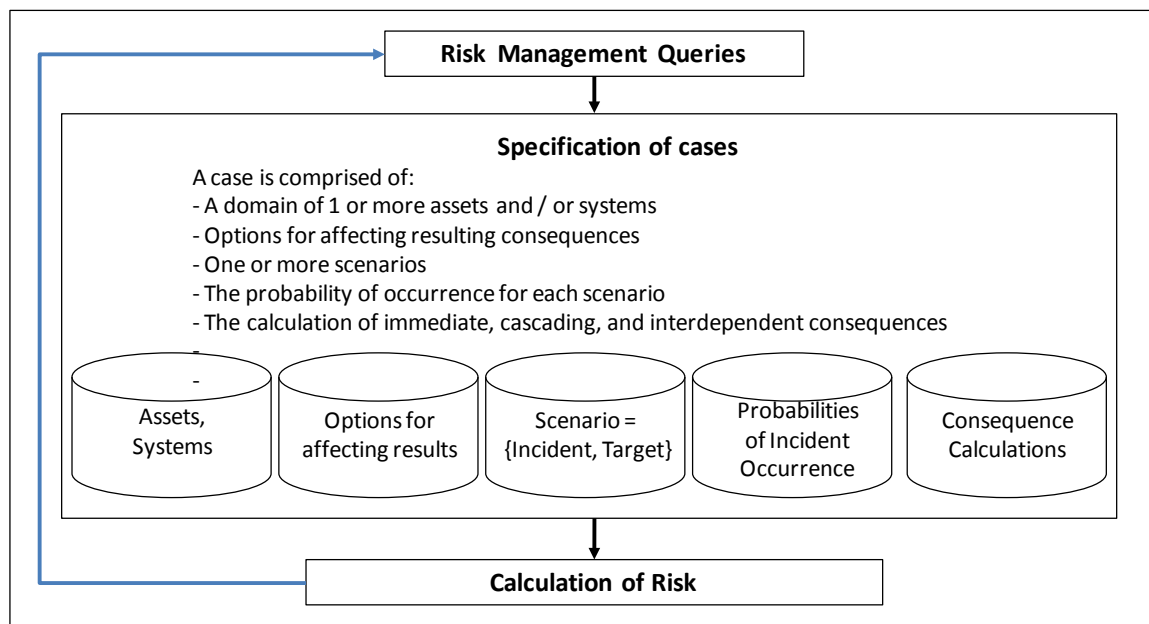


Figure 2. Architecture for Managing a Scenario Space

c. Decision-making Environments and Information Needs

Risk management queries are generated for a variety of objectives and time horizons relative to hypothesized incidents. A sampling of decision-making queries are listed below:

- Evaluating initiatives to improve homeland security (pre-crisis; pre-attack)
 - Policy (e.g., personnel surety programs)
 - Budget (evaluating cost effectiveness of proposed alternatives)
 - Technology investments
 - Enhanced defensive configurations
 - Infrastructure back-up capabilities
 - Stockpiling key resources
 - Research and development (R&D) payoff potential
 - Grant Program
 - Legal/Regulatory
 - Economic incentives to state, local, and private sectors
- Assessing temporary, real-time upgraded security operations (imminent potential crisis)
 - Cost/effectiveness
 - Collateral implications
 - Informing intelligence requirements
- Supporting emergency operations following a terrorist attack (post-attack crisis)
 - Prediction of cascading effects
 - Predicting unintended consequences of operational decisions
 - Real-time operational plan evaluation
 - Informing intelligence requirements
- Supporting clean-up and restoration operational decisions (post-attack, post crisis)
 - Phasing and timing of the restoration of key services
 - Alternatives for key resources
- Defending and supporting budget and policy decisions (pre-crisis; pre-attack)

- Maintaining current risk status of critical infrastructure (all phases—e.g., risk dashboard)
- Building consensus among Federal, state, local, and private stake-holders (all phases)

4. Assessment of Doctrinal Guidelines

A. Cross-Walk with Key Recommendations from National Academies Study

The five doctrinal guidelines address the National Academies study recommendation concerning the formulation of Risk as $R = f(T, V, C)$. Although the study focused primarily on the challenges of developing quantitative estimates of risk for intentional incidents, it also identified advantages in establishing greater commonality and commensurability in the risk management methodologies for both terrorism incidents and natural hazards. This could promote comparative risk analyses and improve the ability to establish priorities across a range of threats and CIKR targets.

The doctrinal guidelines identify several steps to establish greater commonality and commensurability in methodologies:

- Recognize that each risk analysis query has a context
- Use scenarios (threat plus target) as the unit of analysis
- Express quantitative risk in terms of expected value of loss
- Employ multiple consequence metrics
- Use ratio-scale metrics for threat, vulnerability, and certain consequence metrics

The doctrinal guidelines also enable technical solutions to be implemented for the following challenges to DHS risk analysis identified in the National Academies study:³³

- Modeling the decision-making and behaviors of intelligent adversaries
- Characterizing and communicating uncertainty in models, data inputs, and results appropriately
- Resolving methodological issues around implementing risk as a function of threats, vulnerabilities, and consequences
- Modeling cascading risks across infrastructures and sectors
- Incorporating broader social consequences

³³ NRC, *Review*, 51.

- Providing analyses of value to multiple, distributed decision-makers
- Developing risk analysis communication strategies for various stakeholders

B. Conclusion

Following the guidelines in this document will help ensure that concepts and computational methods for estimating vulnerability and the resulting risk calculations produce commensurable risk metrics. Commensurable risk metrics are critical to support cross-sector comparisons of risk. Following these guidelines enables the user to develop transparent, defensible techniques to identify vulnerabilities and assess risks that are asset and systems-based. They provide a way to produce the detailed analytical representations of the assets and systems and their interdependencies while assuring commensurable risk results. Using these guidelines enhances the utility of risk analyses to meet multiple needs, thus saving collaborating partners substantial time and resources.

Appendix A

Estimating Vulnerability for Layered Defenses

A. Overview

This appendix will illustrate how the concept of *layered defenses* and the quantification of vulnerability based on expert elicitation and subjective probability can be the basis for a sound vulnerability doctrine for many critical infrastructure and key resources (CIKR) assets and, in some cases, information technology (IT) systems.

Layered defenses provide a basis for quantifying vulnerability as the joint probability of successfully breaching all relevant defensive layers. This is a doctrinally sound way of estimating vulnerability for many sectors; however, when comparing vulnerability across multiple sectors some of which do not use the layered defense model, it is recommended that the metric for vulnerability be the expected value of loss. For sectors using the layered defense model, this simply means that the probability of successfully breaching all relevant defenses is multiplied by the consequences to produce the expected value of loss metric for comparison purposes with other sectors.

The main document discusses vulnerability from the perspective of how vulnerability information is used by decision-makers. In that discussion, the Institute for Defense Analyses (IDA) study team noted that when large-scale CIKR systems or phenomena are the focus of decision-maker interest, it is common to find sophisticated models and simulations being used to predict the outcome of a terrorist attack or natural event. Moreover, outcomes are often estimated based on loss of lives and economic losses should the event occur. Hence, vulnerability is defined as predicted loss given an event. For example, the vulnerability of a bio-terrorism event would be the predicted loss of lives and loss to the economy should it occur.

To predict the outcome of attacks on simpler CIKR systems and assets, the layered defense model is often a good framework for quantifying vulnerability. The concept of layered defenses, which exists in modern day military tactics and doctrine, is recorded in military annals dating back thousands of years (e.g., the Great Wall of China). More recently, the layered defense model has been applied to cyber security to address the competing dual needs to both protect IT systems while simultaneously using them to share critical information; information sharing—the core function of communications—can, and often does, provide avenues for cyber threats to penetrate cyber defenses to cause both harm and chaos.

The methods documented in this appendix can be used to quantify vulnerability in the context of risk assessments for critical infrastructure assets and systems. Risk assessment is a function of threat, vulnerability, and consequences. In the context of a layered defense model for vulnerability, the following definitions apply:

- Threat—the likelihood of an attack being attempted by an adversary or the likelihood that a hazard will manifest¹
- Vulnerability—the likelihood that an attack will be successful, given that it is attempted²
- Consequence—the estimated outcome of an event, incident, or occurrence on human, economic, mission, psychological and other factors³

Quantifying each of these terms presents unique challenges. Quantifying threat is largely dependent on the availability of information from the intelligence community. Similarly, the valuation of consequences depends on the interdependencies among CIKR assets and systems. Moreover, the damage calculations can be complex and computationally intensive.

IDA's focus is on quantifying vulnerability. The IDA study team took its definition of *vulnerability degree* from the Department of Homeland Security (DHS) *Risk Lexicon*: “the common measurement of vulnerability is the likelihood that an attack is successful, given that it is attempted.”⁴ The designation IDA uses for this is $P(S|A)$ —the probability of success given an attack. This shows that (1) probability theory can be used effectively to quantify vulnerability, and (2) the application of probability theory to the layered defense model can provide an effective basis for rigorously and quantitatively estimating vulnerability in many of the CIKR sectors. When the product of $P(S|A)$ and consequences is determined, the risk metric that results is defined as the expected value of loss from scenarios (potentially many) that can cause loss.

The structure for applying probability theory to the estimation of $P(S|A)$ must be carefully considered and include the following: (1) how layers are identified and defined; (2) how threats to CIKR assets and systems are identified; (3) how scenarios (combinations of threats and targets) are managed; and (4) how consequences are computed. The IDA study team has provided a simple structure that addresses these issues and that can be expanded as needed.

¹ Department of Homeland Security (DHS), Risk Steering Committee, *DHS Risk Lexicon*, 2010 ed. (Washington, DC: Department of Homeland Security), 36.

² Ibid., 38.

³ Ibid., 10.

⁴ Ibid., 39.

The main document provides a set of doctrinal guidelines that enable the user to compute risk in a way that produces commensurate risk metrics across the 18 CIKR sectors. These guidelines are briefly recapped as follows:

1. Define scenarios (combinations of attack vectors/targets) to specify the critical information necessary to estimate scenario risk and its key parameters—consequences, vulnerability, and threat.
2. Ensure that the needs and limitations of the supported decision-making environment are understood and addressed.
3. Use ratio-scale⁵ metrics to quantify threat, vulnerability, and consequences.⁶
4. Quantify vulnerability estimates as an expected value of loss for a given scenario. Two formulations are possible:
 - a. As a product of consequences and probability that an attack is successful—expressed as $C \times P(S|A)$, or
 - b. As a direct estimate of the expected value of loss when $P(S|A)$ estimates are too complex or abstract.
5. Manage the scenario space so that the scope, scale, and assumptions can be verified as appropriate to the decisions, or modified, if determined to be indefensible.

Section B of this appendix provides a detailed example of how to construct a simple layered defense model and apply it to estimate probability of success for physical attacks, such as vehicle-borne improvised explosive devices (VBIEDs) against ground targets. This appendix extends the simple general model to illustrate how it might be applied to small systems like large dams, which consist of many targets that work together to perform multiple functions. Section C further extends the discussion to the more abstract layered defenses of cyber defenses. Section D provides a conclusion to these discussions.

⁵ See Appendix D, “Ordinal Scales and Risk Assessment.” Ordinal scales are commonly used in homeland security risk analyses. Appendix D discusses how they are often misused and provides the justification for using ratio scales. Examples of ratio scale metrics include length, time, probability, and cost.

⁶ Not all consequences can be assessed quantitatively. At the present time, the most common set of consequences that are assessed quantitatively are mortality, morbidity, and economic impacts.

B. Estimating $P(S|A)$ for Simple Physical Defensive Layers—Ground Defensive Layers

1. A Simple Layered Defense Model

Figure A-1 depicts a simple layered defense. The IDA study team uses scenarios to define the threat that causes risk to support risk analyses. A scenario is defined as (1) a specific attack vector (e.g., a sedan carrying a 1,000-lb. improvised explosive device (IED), two attackers, and other information specifying tactics and objectives in cases where needed); and (2) a specific target(s). One or more layers of defense protect the target asset. In Figure A-1, layers A, B, and C protect the target asset. The motivation for using layered defenses is that they are reinforcing—that is, forcing attackers to breach one layer before attacking another (first layer A is breached, then layer B is breached) provides more security than a single composite layer consisting of the combined defensive attributes of A and B (e.g., building blocks of defense such as guards and fences). Layered defenses are sometimes associated with the concept of *defense in depth*—in other words, an attacker is delayed by each layer of defense, allowing more time for the defender to reinforce outer layers, prepare for the coming attack, or mount a counterattack.




Figure A-1. A Simple Layered Defense Model

From a likelihood or probability perspective, the security afforded by a defensive layer is defined as the probability of causing the attack to fail; conversely, vulnerability, is defined as the probability that the attack succeeds. Hence, the vulnerability of layered

defenses is defined as the probability that an attacker successfully breaches all the defensive layers and attacks target asset.

a. Characterizing Defensive Layers for Ground-Based Attacks

The IDA study team characterized the defensive layers that either exist or might be developed for CIKR assets and systems as follows: defensive layers—in most cases—are constructed from available or potentially available defensive measures or attributes. For example, fences, guards, barriers, access control, and surveillance are well-known defensive measures that can protect against ground-based attacks. Table A-1 illustrates how defensive layers are defined by unique sets of defensive attributes.

When defined or characterized by the presence or absence of five defensive attributes, there are total of 32 possible layer defensive configurations (LDCs). Of these 32 LDCs, the 9 judged as most common or feasible with respect to “real world” defensive configurations are shown in Table A-1. Of the potential LDCs not considered, one consisted of only a vehicle barrier without an accompanying fence. This LDC would afford no real protection since the barrier could be easily circumnavigated. Tables such as Table A-1 are extensible, meaning that new defensive attributes and new LDCs can be added as needed.

Table A-1. Layer Defensive Configurations (LDCs) A–O

Attribute	LDC								
	A	B	C	D	E	F	G	H	O
Access Control	✓	✓	✓	✓	✓	✓			
Personnel Barrier	✓	✓	✓	✓	✓	✓	✓	✓	
Vehicle Barrier	✓	✓	✓	✓			✓		
Guard Force	✓	✓							
Surveillance System	✓		✓		✓				

b. Estimating Probability of Success for Individual Layers

Probability estimates (and, especially in this case, combining probabilities for multiple layers) are sometimes subtle and difficult to understand. Probability estimates are made for events that have uncertain outcomes—i.e., given that an attacker attempts to breach a defensive layer in some future hypothetical scenario, it is not known with

certainty whether the attack will succeed. Probability theory is based on the premise that the outcomes—success and failure—must be clearly defined.⁷ Hence, it is important to define the criteria for “success.” For this document, success is defined as breaching a defensive layer with enough attack vector resources remaining to accomplish damage (or destruction) that meets or exceeds an established measure of consequence. Analysts or decision-makers applying the layered defense model must establish the criteria for damage/consequence to the target for each attack vector/target scenario.

One straightforward approach for developing Probability of Success estimates for each LDC is to use experts and expert elicitation. In addition to well-defined LDCs, experts also need well-defined threat attack vectors. For example, ground attack vectors can be considered as a class of attacks against ground LDCs. Well-defined instances of ground attacks vectors might include a single attacker with a backpack filled with explosives, a sedan filled with 500 pounds of explosives, or a semi-truck/trailer filled with 30,000 pounds of explosives. Identifying attack vectors and matching them with targets to create scenarios are important and continuing management challenges. Threat information that defines the potential attack vector space often originates with the intelligence community and can be combined with historical cases of real attacks. The overarching criteria for managing the threat space using defined, representative attack vectors are to ensure that all potential threats that cause risk are included and have representative attack vectors, and that redundancies are eliminated.

Once the threat has been characterized by specific, well defined attack vectors and defenses characterized by specific LDCs are identified (Table A-1), expert elicitation using qualified and experienced experts can be used to estimate probability of success for each layer, as shown in the cells of Table A-2.

⁷ Nozer D. Singpurwalla, *Reliability and Risk* (Chichester, England: John Wiley and Sons Ltd., 2006).

Table A-2. $P(S|A)$ Estimates for Each Attack Vector–LDC Combination

Attack Vector	LDC								
	A	B	C	D	E	F	G	H	O
AV-1									
AV-2									
.	<i>Each cell contains an estimate of $P(S A)$ for the attack vector / LDC combination shown</i>								
.									
.									
AV-M									

c. Combining $P(S|A)$ Estimates for Multiple Layers

The overall objective is to estimate the probability of success for all layers combined for a given target asset. If the probability estimates for each layer are assumed to be independent,⁸ the probability of penetrating all defensive layers is the simple product of the individual layer probabilities. For different attack vectors, this product can change, since the $P(S|A)$ estimate for each LDC shown in Table A-3 depends on the attack vector.

Figure A-2 illustrates how the overall $P(S|A)$ estimates are combined. For illustration purposes, IDA assumes that layer A is LDC “C”; layer B is LDC “F”; and layer C is LDC “H.”

⁸ Conditional probability estimates provide a more general case; however, more features and considerations must be added to account for dependencies among the layers and their effect on the estimation of conditional probabilities.

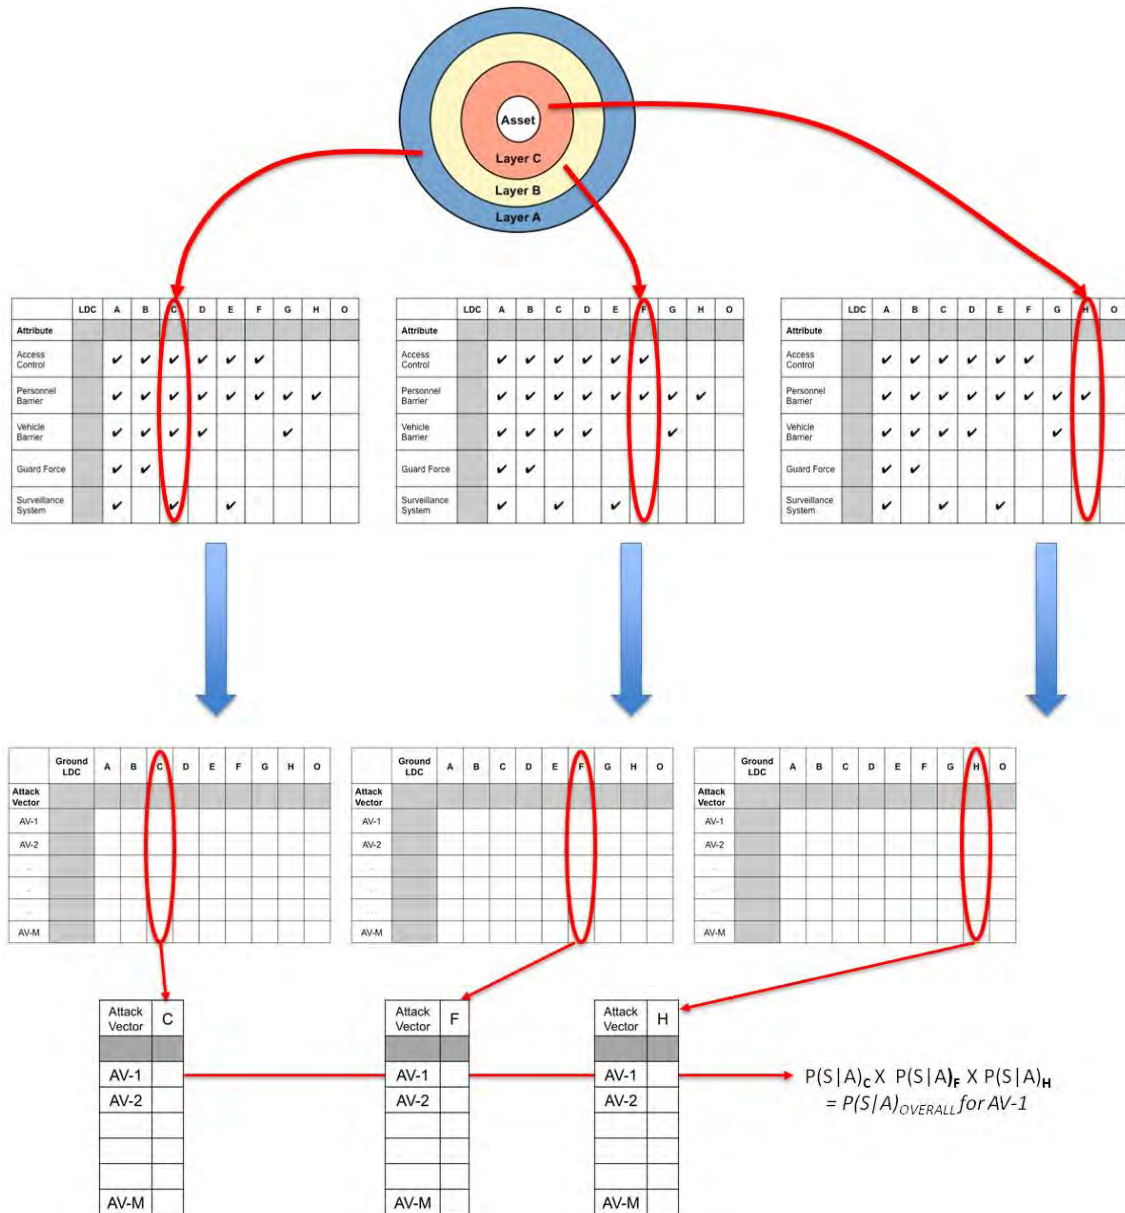


Figure A-2. Illustration of Using Looking Up Tables for Probabilities of Success Generated by Expert Elicitation and Combining Them for the Overall Joint Probability of Success

2. More Complex Defensive Layers

IDA expanded the concept of simple layered defenses by using dam projects. Dams represent a class of targets that can contain multiple high value targets. For example, a given dam project might have the following assets: the main impoundment structure; spillway gates for controlling water pool level; hydroelectric generation; and locks to support commercial shipping. Although each of these assets can have its own combined layered defenses, the approach for estimating $P(S|A)$ is the same as shown in Figure A-3.

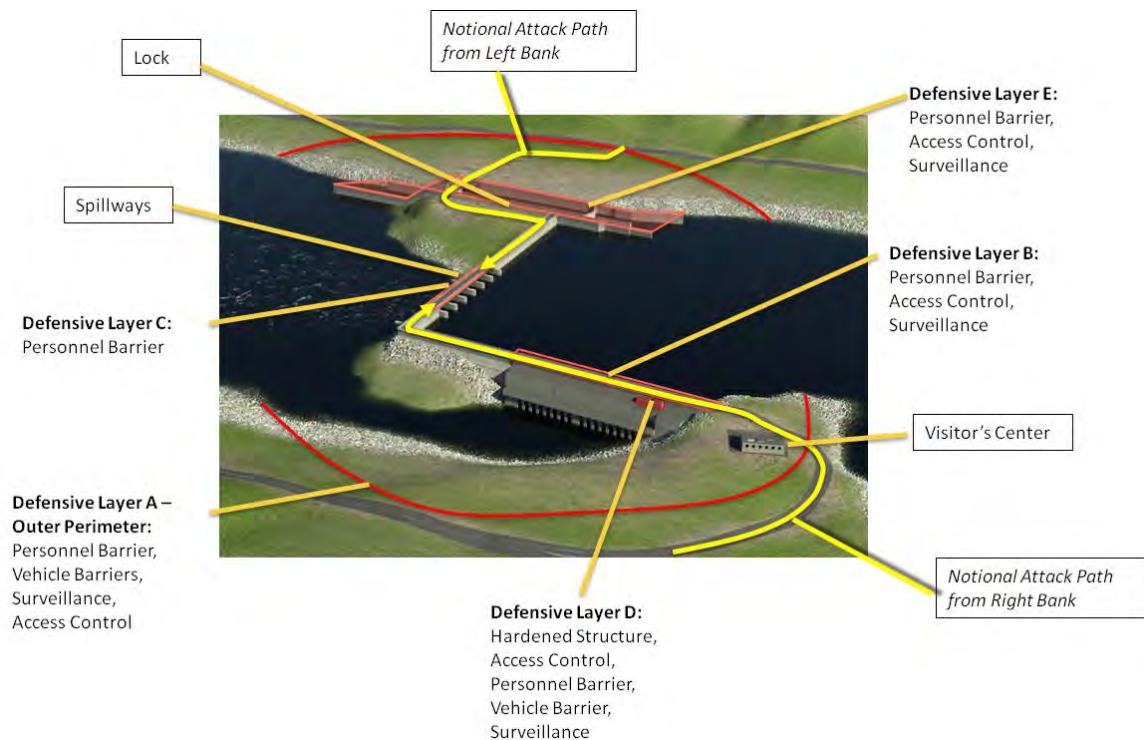


Figure A-3. Notional Graphic Showing Layered Defense of Dams

As Figure A-3 shows, dams offer an additional layer of complexity because attacks can originate from the left or right bank. As a result, the same attack vector approached from opposite sides can have different overall combined LDCs since the sequence and characteristics of layers may be asymmetrical. Figure A-3 also shows threat vector approaches from the water—both upstream and downstream. The construction of P(S|A) tables for water-borne attacks would proceed the same way that ground attack P(S|A) tables were explained in Section B.1.

C. Estimating P(S|A) for Abstract Layers—Cyber Layered Defenses

The control functions of installations like dam projects may be through an automated, computer-based digital control system—supervisory control and data acquisition (SCADA) system. SCADA systems are a specific type of automated process control systems (PCS). One or more of these systems may exist at a dam project. SCADA systems may be connected to a larger network. In keeping with the layered defense model described in Section B for estimating vulnerability of targets to ground-based attacks, the vulnerability of SCADA systems can be characterized by: (1) the targets—the PCS and the project assets/functions it controls; and (2) the defensive layers—described using the cyber defensive configurations (CDC) in place to protect it. A CDC is analogous to the LDCs that defend against ground attacks or water-borne attacks.

The purpose of this section is to illustrate a notional construction of CDCs for each dam computer-based control system. The approach for doing this is similar to LDC/attack vector construction for ground-based scenarios, but with a different set of defensive attributes that, when combined, can defeat a range of cyber threats. First, the study team identifies the attributes, then it constructs real world, feasible CDCs, then it develops a table of P(S|A) estimates for a selected set of representative cyber attack vectors. The cyber defenses are independent of the physical defensive layers because a cyber attack can be conducted from a remote location via an Internet connection, or by a trusted insider who has access to a computer within the dam project.

1. Characterizing Defensive Layers for Cyber Based Attacks

Most PCS IT systems can be represented by three separate configurations based on the increased automation sought. These configurations may or may not have been developed with security considerations in mind, yet they all have an impact on security. All of these configurations will exist in the vast array of critical infrastructure. Although this appendix describes three types of PCS, integrated, separated, isolated, it provides detail on only one of these configurations to illustrate how to construct a layered cyber defense model for dam SCADA systems. The three types of PCS systems are illustrated in Figures A-4 through A-6.

The following list provides a short description of the components for each PCS.

- **Web Server**—a computer that hosts web sites to deliver accessible content via the Internet. It includes the hardware, operating system, Web server software, transmission control protocol/internet protocol (TCP/IP) protocols, and site content (Web pages, images and other files)
- **Administrative (Admin) PC**—a computer used as an administrator's console
- **Mail Server**—a computer that serves as a virtual post office by storing incoming mail for distribution to users (message store) and forwarding outgoing mail through the appropriate channel (message transfer agent).
- **Marketing PC**—a computer used in marketing and sales with access to current production and use
- **Program Logic Controller**—a computer that accepts commands from the PCS Master and executes process control
- **PCS Historian**—a computer that keeps historical sales, usage and outage data for use by sales and control people
- **PCS Master**—a computer that controls the process control elements, accepts inputs from the admin and process people
- **Admin Historian**—a computer that stores past data for prognostication and planning.

a. Integrated

In an integrated environment, the PCS is one more component of a corporate network. All components are physically tied together through an Ethernet cable or other networking approach. From a corporate standpoint, this is the most convenient and easiest to implement, but the most vulnerable. Examples of a PCS may be a small chemical plant, food processing plant, or a single site where all aspects of the service are provided, such as a small rural electric company.

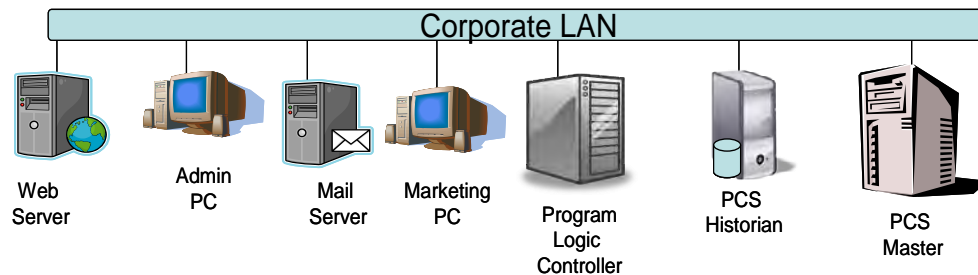


Figure A-4. Integrated Process Control System

b. Separated

In a separated environment, it is recognized that the PCS needs to be separated from the overall network for increased security. This is often done to prevent corporate or marketing from directly affecting the PCS, which may have very stringent real-time requirements. In order to do this, there is an admin historian set up with a refresh rate that is consistent with corporate needs and that provides a near real-time picture of the operation. Although they are separate networks, they are logically connected through the update of the admin historian, which is used by the corporate and marketing sides of the house. This also has the added benefit of improving the overall security of the PCS. A separated PCS might be found in an intermediate-size water facility or electric power operation.




Figure A-5. Separated Process Control System

c. Isolated

The third type of configuration is an isolation approach, which may be the result of safety configuration concerns or legal requirements because the risk of being connected to the corporate network is too great. In this configuration, the logical link is broken and a one-way update to the admin historian is undertaken. This can be accomplished in several ways from an IT perspective, but one way is to post a message on a “blackboard” from the PCS side and read it from the blackboard on the corporate side.⁹ This increases the lag time in communications with the corporate network, but it can still be very near real-time. One example of this arrangement might be a dam SCADA system.




Figure A-6. Isolated Process Control System

2. Cyber Defensive Attributes for Cyber Defensive Configuration

For each of the different configurations—integrated, separated, and isolated—the next step is to develop a set of cyber defensive attributes that can be used to characterize CDCs. The following is a list of defensive attributes and definitions for the CDC.

- **Firewall and anti-virus/spam**—the first line of defense and probably the easiest for a corporation to implement. It implies that some form of identification and authorization are in place. These may range from simple passwords to biological identification. The strength of function associated with firewalls and identification are assumed to escalate with the overall security posture.
- **DMZ for Internet and intrusion detection**—considerably more difficult to implement, requiring more hardware and sophistication. The DMZ is essentially a set of firewalls that act as one-way valves to keep hazardous activities contained where they can be monitored and prevented from spreading. The name “DMZ” is derived from the military term “demilitarized zone”; it is also referred to as a perimeter network.

⁹ A blackboard is an isolated partition of memory that only permits the PCS Historian to write to it. All other access is “read-only”—for example the Admin Historian has “read only” access to the blackboard.

- **Intrusion Detection**—hardware and software dedicated to checking the information flow for anomalous behavior that might indicate a cyber intrusion. The simplest form of this is the *honeypot*, which consists of a work station on the net that contains files that are attractive looking to an intruder and software that notifies a security individual when these files are accessed. Under ordinary circumstances, the honeypot is not accessed by trusted or untrusted personnel in the IT environment. An important concept here is that indications of intrusion are met by a response team that will undertake the forensics and cleaning of systems. Many attacks are discoverable in the early phases before serious damage can occur.
- **Actively managed security** (training, audits, updates, log reviews)—a change in corporate philosophy making security every employee’s responsibility. This attribute is not only essential to a secure computing environment, but it also bolsters the response-team approach described in intrusion detection.
- **Defense in depth**—additional firewalls and intrusion detection at key points, including PCS encryption and authentication. This attribute adds extra impediments to an attacker by making him overcome additional obstacles. It also provides a point of awareness for intrusion detection systems.
- **Insider Protection** (personnel screening and monitoring) —a measured amount of paranoia is required at this level. The insider threat is the most difficult to stop.

Table A-3 shows five defensive attributes and an associated set of six CDCs.

Table A-3. Cyber Defensive Configurations (CDCs) Based on Unique Sets of Cyber Defense Attributes

Cyber Defensive Attributes	CDC					
	1	2	3	4	5	6
Firewall, anti-virus protection		✓	✓	✓	✓	✓
DMZ, intrusion detection			✓	✓	✓	✓
Actively managed security				✓	✓	✓
Defense in depth, VPNs					✓	✓
Personnel checks, insider protection						✓

3. Estimating Probability of Success for Cyber Defensive Layers

As with ground-based attacks and LDCs, attack vectors (cyber attack vectors—C-AV) must be defined for cyber defenses in order to provide cyber experts with the information needed to make informed probability of success estimates. Cyber attacks can be more complex and evolve at a much more rapid pace than ground-based attack vectors. Hence, the management of the scenario space is even more critical in the cyber threat arena than in a conventional attack context. Change occurs so rapidly that at times cyber defenders are forced to react to the more recent attack rather than plan and implement defenses over a long period of time, as is the case for physical defenses (e.g., the construction of fences or barriers). Nevertheless, the standard for defining cyber attack scenarios is the same: sufficient information must be given to cyber experts to be able to estimate the probability of success for the cyber attack vector against the cyber defenses.

The IDA study team provided three cyber attack scenarios to illustrate a layered defense protection of PCS against attacks:

- A cyber-terrorist group is assumed to be well funded, have extensive expertise in network and computer technologies, and have adequate time for network surveillance and attack planning. This group would have the resources and time to set up a model network that emulates their target environment in a laboratory setting, where they could practice attack techniques. They are also likely to have the time and resources to plant one or more members of their group as employees inside the target organization.
- Individuals with moderate cyber capabilities and the advantage of time and timing, have the tools available, including those commonly available on the Internet and through *Black-Hat* contacts.¹⁰ The extortion group's motivation and objective is money, not disrupting the victim's operations. The group will threaten cyber attack to disrupt or shut down control system capabilities unless a sizeable ransom is paid. The group would be just as happy to be paid and not attack the target network. If it is not paid, it has a fair amount of expertise to penetrate networks and disrupt computer systems. It will take care not to be identified and to ensure that the attack cannot be traced back to the source. The objective is to teach the victim a lesson so that the group's extortion demands will be paid the next time it threatens. It is perfectly acceptable to the extortion group if the attack looks to the public like an accidental disruption of service. This threat is typified by a criminal extortion, but may also be perpetrated by a disgruntled or greedy employee.

¹⁰ Black-Hat organizations are groups of hackers that share information through electronic means and through annual meetings, typically in Las Vegas.

- A small group of individuals with moderate cyber capabilities and a limited timeframe to execute a plan, either through coordination with other events or impatience. This group has a fair amount of expertise to penetrate networks and disrupt computer systems. The tools available include those commonly available on the Internet and through Black-Hat contacts. The motivation is to disrupt or shut down electrical power or oil or gas distribution as a political, social, or religious statement. The group may also include disgruntled employees. The objective is to cause damage to the particular utility and inflict injury on the population it serves. These groups are more likely to time their attacks to coincide with bad weather, such as a winter cold spell to expand the effects of a power outage. They are less concerned about being identified; in fact, some may openly claim responsibility. The group is expected to disband immediately after a successful attack. This scenario is typified by a spot terrorist attack, but could be any of the groups mentioned above.

Once the threat has been characterized by specific, well-defined attack vectors, and defenses have been characterized by specific, well-defined, layered CDCs (including integrated, separated, and isolated PCS), expert elicitation using qualified and experienced experts can be used to place probability estimates for each cyber attack scenario into the cells of Table A-4 against each cyber attack scenario.

Table A-4. P(S|A) Estimates for Each Attack Vector–CDC Combination

Cyber Attack Vectors	CDC 1	CDC 2	CDC 3	CDC 4	CDC 5	CDC 6
	Integrated					
C-AV-1						
C-AV-1	<i>Each cell contains an estimate of P(S A) for the attack vector/CDC combination shown</i>					
C-AV-2						
C-AV-3						
C-AV-4						
C-AV-N						

D. Conclusion

The discussion in this appendix has focused on developing a layered defense model for quantitatively estimating the vulnerability metric for the DHS risk metric. Risk is a function of threat, vulnerability, and consequences. Vulnerability in the layered defense

model examples presented above is defined as the probability of success for a given attack vector in breaching all defensive layers that protect the intended target. The IDA study team defines a scenario as the combination of an attack vector and the intended CIKR target asset.

Three examples were provided, each showing the steps to construct a layered defense model for representative CIKR targets: (1) a simple layered defense consisting of concentric layers of defense against ground-based IED attacks; (2) a more complex layered defense protecting a small CIKR system (dams); and (3) a more abstract layered defense protecting process control systems against cyber attacks. These examples illustrate the flexibility of the conceptual layered defense model to accommodate diverse applications for estimating vulnerability of CIKR assets and systems.

As the main document discusses, there are two ways to interpret vulnerability. The layered defense model presented in this appendix follows the interpretation of vulnerability as a probability of attacker success. This interpretation works especially well with the layered defense approach and, as illustrated, can be used for complex systems as well as assets. This interpretation is well suited when the decision environment is focused on improvements to the system of layered defenses. In other cases, the layered defense model is difficult to develop—especially when a CIKR system is very complex and does not lend itself to a layered defense characterization. In these cases, defining vulnerability as the expected value of loss can be more appropriate, as is shown in Appendix B which explores biological attack defenses. Even in bio-terrorism example, however, both models may be appropriate for use. For example, inoculations against certain bio-pathogens can be abstractly viewed as a defensive layer and there may be other defensive layers that could be similarly used. It is up to the decision-maker to decide which approach works best for the issues at hand.

Appendix B

Estimating Vulnerability for Biological Attacks on U.S. Populations

A. Introduction

The work documented in this appendix is part of a larger study by IDA to prescribe a procedure for the quantification and consistent application of risk assessment methods. The purpose of this appendix is to examine the process for risk assessment as it applies to a contagious biological agent event.

This appendix outlines the process for assessing risk associated with contagious biological agent events as a function of the number of lives lost by providing a review of the biological agent scenario from the strategic perspective; reviewing the guidelines with a focus on a contagious biological agent event, and then providing a very simple example.

1. Background

Biological diseases and the potential for epidemics and pandemics, both naturally occurring and accidentally or intentionally spread, pose several challenges for public health and homeland security. Naturally occurring epidemics have decimated populations locally, regionally, and even globally; the Black Plague in the sixteenth century is believed to have killed approximately a third of the European population, while the Spanish Flu epidemic (1918–1919) spread globally, killing more than 50 million people. Historically, diseased blood, fecal matter, and bodies have been used to create toxic artillery and poison water supplies. In addition, several countries are known or alleged to have developed weaponized biological agents. Contagious biological agents provide some of the earliest documented uses of chemical, biological, radiological, and nuclear (CBRN) agents in warfare and remain a concern today, as potential weapons of both governments and terrorists.

Contagious biological agents, either naturally occurring or intentionally spread, may result in a surge of cases flooding the public health sector and draining available resources. Contagious biological diseases may require introducing unique policies and change the dynamics of social interactions as “shelter-in-place” and “mandatory-personnel-only” orders are given and the public self-selects social distancing (the voluntary reduction of daily social interactions among the general, susceptible population) to limit potential interactions that could lead to contracting the disease. Likewise, mandatory herd-vaccinations, medical countermeasure distribution procedures, and contact tracing may be implemented to minimize the spread of disease.

Understanding the potential for disease spread as well as the implications of the wide variety of political, educational, and medical countermeasures that can be implemented before an event to help mitigate disease spread, post-contagious biological agent introduction, and initial infections is extremely important for the planner and policy maker.

The threat posed by contagious biological agents, which is an evolutionary event continuously progressing and changing, is similar to those posed by a single-point hazard (i.e., an improvised explosive device detonation) but also slightly different. In one sense, they are similar: as with a single-point hazard, layers of defense can be put into place in advance. Technological detectors and identifiers can be used much like x-ray scanners to detect the presence of agent. Similarly, physical measures, including standoff distances and filtered ventilation systems, may provide some protections. Additional layers of defense may be applied at the individual level—for example, vaccination and prophylactic medicine (post-exposure, pre-symptom onset).

Unlike a single-point hazard, however, there is no clear moment that the contagious biological event starts; the event may be considered to start at a number of different points depending on the questions being asked: (1) the introduction of the agent (e.g., intentionally, when released from a sprayer vice naturally, when the gene mutates to become more infectious and is introduced into a human host); (2) when the first infection begins, which may not be identifiable until after the first patient has presented and/or died; and 3) when the number of infections reaches some pre-established threshold.

Further, because a contagious biological event is constantly evolving and changing, the introduction of new policy measures, medicines, education, and procedures will be preemptive for those who have not yet been exposed and responsive for those who have. Thus, contagious biological agents pose a unique challenge when assessing risk and vulnerability.

2. Risk Assessment and Vulnerability

Risk assessment is a function of threat, vulnerability, and consequences. For the purposes of this appendix, the following definitions apply:

- Threat is the likelihood of an attack being attempted by an adversary or the likelihood that a hazard will manifest.
- Vulnerability is the expected value of loss, given that the attack is attempted.¹
- Consequence is the effect—human, economic, missions, and psychological, as well as other factors—of an event, incident, or occurrence.²

¹ Another definition for vulnerability is the “probability that an attack will be successful, given that it is attempted.” For biological threats and other threats that pose risks to systems, sophisticated computational models are capable of quantifying losses directly, and these results are often a more meaningful and rigorous estimate of vulnerability for decision-makers.

The quantification of threat is a challenge and is largely dependent on the availability of information from the intelligence community. Similarly, the valuation of consequences is challenging and depends on the physical damage that actually occurred during an event.

For the purposes of this risk assessment, the risk metric is defined as the expected value of loss from—potentially many—scenarios that can cause loss. In essence, this is a quantification of vulnerability. To quantitatively assess vulnerability, the following guidelines apply:

1. Define scenarios (combinations of attack vectors/targets) to specify the critical information necessary to estimate scenario risk and its key parameters—consequences, vulnerability, and threat.
2. Ensure that the needs and limitations of the supported decision-making environment are understood and addressed.
3. Use ratio-scale³ metrics to quantify threat, vulnerability, and consequences.⁴
4. Quantify vulnerability estimates as an expected value of loss. Two formulations are possible:
 - a. As a product of consequences and probability that an attack is successful, or
 - b. As a direct estimate of the expected value of loss when $P(S|A)$ ⁵ estimates are too complex or abstract.
5. Manage the scenario space so that the scope, scale, and assumptions can be verified as appropriate to the decisions, or modified if determined to be indefensible.

Further, to understand the risk being assessed and the vulnerability being quantified, the loss itself must be defined. Loss is often expressed in terms of lives—lives lost and/or time lost due to illness or injury—and monetary loss—including actual economic loss, property damage and destruction costs, restoration costs, and lost wages, among other monetary losses. Depending on the threat, other losses may be considered; however, in order to establish a metric for comparison across a wide variety of threats and because following a biological event, the most

² Department of Homeland Security (DHS), Risk Steering Committee, *DHS Risk Lexicon*, 2010 ed, (Washington DC: Department of Homeland Security, 2010).

³ See Appendix D, “Ordinal Scales and Risk Assessment.” Ordinal scales are commonly used in homeland security risk analyses. Appendix D discusses how they are often misused and provides the justification for using ratio scales.

⁴ Not all consequences can be assessed quantitatively. Currently, the most common set of consequences to assess quantitatively includes: mortality, morbidity, and economic impacts. Analysts can qualitatively (e.g., using ordinal scales) describe the associated psychological/behavioral consequences and the impacts on missions, and keep separate the estimates for vulnerability, which can be quantitative or qualitative if they are not numerically combined with consequences.

⁵ $P(S|A)$ is the conditional probability of success given the attack A (noted by the vertical bar, “|”); it summarizes the essential notion of vulnerability to a threat.

common questions will focus on the human element of the event, this appendix will focus specifically on lives lost, lives injured, and, potentially on the monetary valuation, as feasible. Immediate questions include: How many lives were lost? How many are being hospitalized? For how long? Later, the questions will likely turn to economic concerns.

3. Objective

It is important to recognize the challenges associated with modeling a contagious biological event. Because the agent is contagious, there is the likelihood that the disease will spread regionally, nationally, and even globally. From a systems point of view, the global, interconnected sociological system has the potential to be severely harmed by either a naturally occurring or man-made biological event. What starts as a local event can quickly escalate into an epidemic or even a pandemic before the first case is even identified and diagnosed. For example, because of both the incubation periods associated with contagious biological agents and the speed and rate of travel, an individual exposed in one city may travel around the world, potentially exposing others, and eventually manifest the disease in a city thousands of miles away from the initial exposure point. Multiplied by hundreds of initial exposures, the end result is cities across the globe simultaneously experiencing outbreaks of contagious disease. The potential exists, from the modeling perspective, for the problem to become so large so quickly that it is overwhelming for the modeler.

The objective of this appendix is to illustrate how the definitions and guidelines above can be used to assess and estimate vulnerability and risk from a biological threat due to both naturally occurring and man-made biological agents. The intent is to show how both risk and vulnerability estimated for a complex biological event can be formed in such a way that the results are commensurable with risk and vulnerability estimates computed in critical infrastructure sectors using the definitions and guidelines presented in the previous section. While some sectors may calculate localized risk and vulnerability, biological events, like other cascading events, involve risk to multiple systems and critical infrastructures. The IDA study team intends to demonstrate that, while the methods of estimation can vary significantly across sectors, the same desired outcome of commensurability may still be achieved.

B. Contagious Biological Agents from a Strategic Perspective

As already noted, contagious biological agents pose a unique threat for a number of reasons including, but not limited to, the low likelihood of detection given a biological attack, the duration between exposure and symptom manifestation, the possible dispersion of initial exposure during the incubation period, and the potential for transmission from those initially exposed to close contacts and beyond. Contagious biological agent events include those caused by naturally occurring biological agents, such as influenza and severe acute respiratory syndrome (SARS), and those resulting from intentional weaponization and release of biological warfare agents, such as smallpox and plague.

Whether the event is naturally occurring or intentionally initiated, the initial exposure to a biological agent will result in some fraction of the population becoming exposed and infected. It is generally assumed that individuals must exhibit clinical symptoms and signs for a contagious biological agent disease to be transmissible to others. Following exposure, there will be an incubation period before symptoms and signs manifest,⁶ after which the exposed and infected population will begin to manifest symptoms. The symptoms themselves may, initially, confound diagnosis; symptoms, especially in the early stages of the disease. The symptoms that accompany the initial stages of many contagious biological agents are described as “flu-like” and include such non-specific symptoms as headache, fever, malaise, and muscle ache or weakness. A complete diagnosis for the first several index cases may only be possible after the first specific symptoms manifest.⁷

In the meantime, the biological agent exposure may or may not have been detected by a Biowatch or similar sensor. Confirmation of an exposure, however, will likely be the first positively diagnosed case to surface multiple days after the exposure event or the beginning of the outbreak. It may take several diagnosed cases before an epidemic or the exposure location is identified and interventions are introduced. Several interventions, including medical interventions and policy-driven interventions, will then be considered at the local level and up through national and international government levels.

Medical interventions include post-exposure prophylaxis and antibiotic/antiviral initiation. It is worth noting that post-exposure prophylaxis, or vaccination, will likely be administered to many people who have not yet been exposed to the disease; post-exposure refers to the time in the general spectrum of the disease, after the initial exposure, when some portion of the population has been exposed to the disease. Vaccines exist for several contagious biological

⁶ For contagious biological agents, and indeed most biological agents, this incubation period poses the biggest challenge to public health since exposed individuals travel away from the source of exposure, as far as across the country or around the world, before the illness manifests and the source of exposure is identified. This results in multiple simultaneous index cases in often unrelated and decentralized locations. Until the source of infection is identified and information regarding the illness is spread to global medical resources, each infected, symptomatic individual has to seek medical treatment while possibly transmitting the disease to others and be diagnosed and treated by doctors who may or may not be familiar with the biological agent or disease.

For the purposes of this appendix, the IDA study team, therefore, considers and discusses a localized infection with possible diagnosis information shared through the medical community, and the ability to implement medical and policy countermeasures. The methodology described could be expanded to consider a larger infected space, however, identification and subsequent countermeasure initiation become more complex the larger the area considered and the larger the number of distributed initial infections modeled.

⁷ Specific, identifying symptoms may include bloody sputum for plague patients or the maculopapular rash associated with smallpox. While laboratory diagnostics exist to identify these diseases, the early presentation of symptoms may not prompt testing for such uncommon diseases without additional information or environmental detection triggers. Lien-Teh Wu, 1926, *A Treatise on Pneumonic Plague*, C.H.474, Geneva: League of Nations Health Organization; Donald A. Henderson et al., 1999, Smallpox as a Biological Weapon: Medical and Public Health Management, *Journal of the American Medical Association (JAMA)* 281(22): 2127–37; and David R. Franz et al., 1997, Clinical Recognition and Management of Patients Exposed to Biological Warfare Agents, *JAMA* 278(5): 399–411.

agents. Some, like the influenza, diphtheria, and measles vaccines, are given routinely to the majority of the population; others, like the smallpox vaccine, are only routinely administered to a small portion of the population, including emergency responders and health care workers. Larger vaccination clinics or medical distribution centers would be necessary to provide large-scale population protection.

Whether policy-driven or social, interventions may include, but are not limited to, social distancing, quarantining of the likely exposed, and patient isolation. Shelter-in-place recommendations and social distancing may be initiated by the population itself or through government recommendations, including closing schools and suspending religious and other community meetings. Quarantining the potentially-exposed involves removing individuals who have likely been exposed and are expected to manifest the disease from the population before the onset of symptoms and possible transmission; this may involve home-quarantine or quarantine in a specifically-identified location. Similarly, isolating patients who have already manifested symptoms in their homes, in established medical care facilities or in alternate locations, is intended to reduce contact transmission by removing the patients from the population they could potentially infect.

Each of these interventions requires the establishment of policies, planning, and education. Therefore, the ability to measure and quantify each intervention's ability to reduce vulnerability and compare potential interventions and combinations has potential value for the planner.

1. Schematic

The following schematic depicts a notional contagious biological agent human response estimation process and demonstrates how just a few of the interventions mentioned above might affect different segments of the population during an epidemic. The block diagram at the top is the susceptible, exposed and infected, infectious, removed, and prophylaxis efficacious (SEIRP) model⁸ for estimating contagious biological response,⁹ developed as part of the Human Response Injury Profile (HRIP) methodology for casualty estimation. The SEIRP model estimates contagious biological agent human response as a function of time and time-varying transmission rates. This model is discussed in more detail below.

Contagious biological agent human response is modeled as a progression of individuals through cohorts describing the stages of injury. Prior to exposure, the population is susceptible to infection. Following exposure, some fraction of the population is exposed and infected and has received a dose sufficient to develop symptoms of the disease following some incubation period. Each day, some fraction of the exposed and infected population develops symptoms and signs

⁸ D. S. Disraelly, T. J. Walsh, and C. A. Curling, "A New Methodology for Estimating Contagious Biological Agent Casualties as a Function of Time," *Mathematical and Computer Modeling* 54 (2011): 648–659.

⁹ Other human response and casualty estimation models could be used. The authors are familiar with this model and chose it for illustrative purposes.

and is, therefore, considered ill and/or infectious. After the illness has run its course, fractions of the population are removed from the medical system, either because these individuals have become fatalities or because they have recovered.¹⁰

A notional timeline of the epidemic is shown in Figure B-1, along with some of the mitigation and response activities that could be undertaken. The arrows indicate which fractions of the SEIRP population may be impacted by the introduction of the mitigation and response activities.

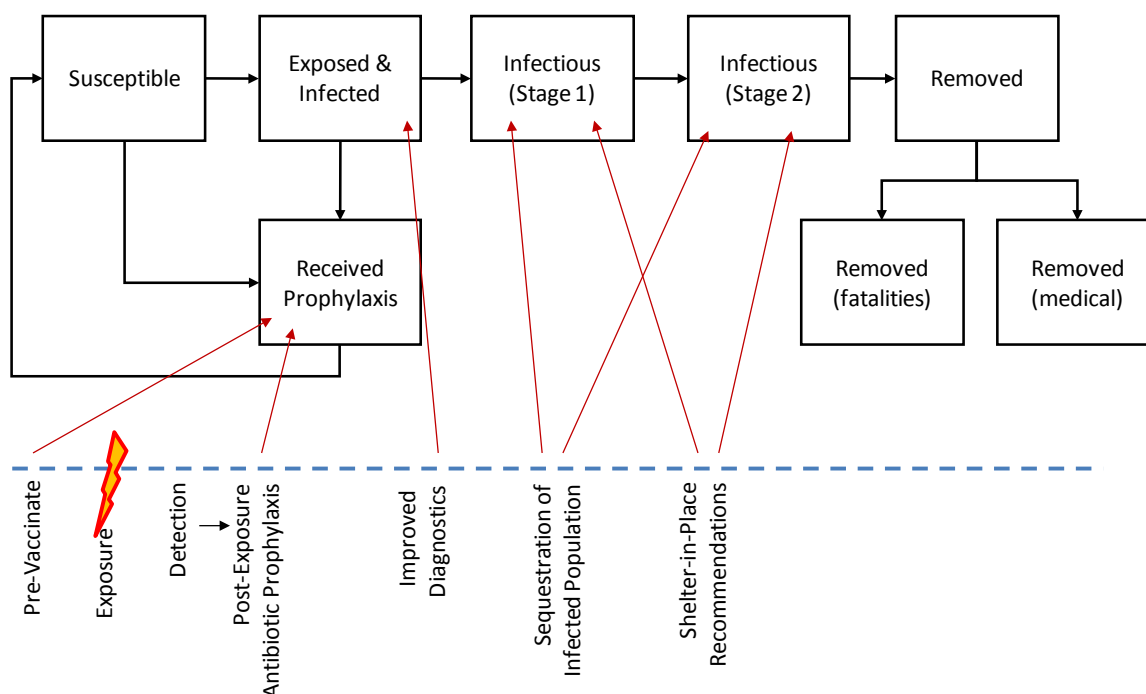


Figure B-1. Notional Schematic of a Contagious Biological Agent Defense

In this example, there is one risk mitigation intervention that could potentially be undertaken prior to the event. The application of pre-vaccination to some portion of the population—e.g., first responders and first receivers—or to the whole population would change the fraction of the population for whom prophylaxis is efficacious before the event occurs.

After the event occurs and the epidemic begins, improved diagnostics could lead to a higher number of people being treated early and therefore not developing the disease and not becoming infectious towards others. Likewise, sequestering the infectious population or limiting potential interactions, and, therefore, the potential for transmission, through shelter-in-place recommendations and social distancing, limits the potential for agent spread, thereby reducing the transmissivity of the disease from the infectious cohort.

¹⁰ Disraelly, Walsh, and Curling, “A New Methodology,” 648–659.

2. HRIP Contagious Biological Agent Methodology

NOTE: The HRIP methodology, using SEIRP,¹¹ is described for the purposes of this example only. Several possible models or even another accepted set of guidelines and rules could be used to estimate vulnerability as a function.

The HRIP methodology was developed to provide planners with an improved ability to estimate casualties resulting from exposure to CBRN agents or effects. In this methodology, to estimate vulnerability a human response model “determines the effects (expected value of human loss) of CBRN exposures on individuals/populations. It calculates the type and severity of illness or injury suffered by individuals, as well as their subsequent death or recovery.”¹² The HRIP methodology relies on time-based progressions of underlying symptoms and signs and their severity changes to determine user-defined casualty estimates following CBRN events.

The HRIP methodology provides the capability to estimate the time-dependence of the numbers of individuals wounded and killed, which aids in the development of medical and logistics requirements, acquisition strategies, and planning. The methodology also characterizes the incidence of injury by severity to provide additional information to planners. Further, the HRIP methodology allows the user to define the casualty threshold—the severity at which casualties occur; this means that the user can determine patient loads and how they might change as a function of injury (and sign and symptom) severity. Decision-makers can use this information to mitigate the vulnerability and risk associated with a biological event.

The HRIP contagious biological agent methodology is built on the SEIRP human response calculation with a time-varying transmission rate. In its simplest form, as illustrated in Figure B-1, the SEIRP model employs time-varying cohorts to describe the dynamics of an epidemic. A series of time-dependent difference equations are solved to estimate human response over time.

As noted above, other methodologies exist for estimating casualties resulting from contagious biological agent exposure. A risk analysis could use any of these other methodologies to calculate the expected value of human loss. The IDA study team selected the HRIP methodology, however, because of its representation of illness and contagion spread over time, its effectiveness as a modeling tool, and the ease in calculating vulnerability in terms of lives lost and resources needed as a function of the disease and the interventions implemented.

3. Typical Queries

As discussed earlier, vulnerability to a contagious biological event or epidemic is measured, first and foremost, in terms of the expected number of lives lost. In contrast, the efficacy of

¹¹ Disraelly, Walsh, and Curling, “A New Methodology,” 648–659.

¹² North Atlantic Treaty Organization (NATO), 2010, *Allied Medical Publication 8(C)—North Atlantic Treaty Organization (NATO) Planning Guide for the Estimation of Chemical, Biological, Radiological, and Nuclear (CBRN) Casualties*, Ratification Draft (AMedP-8(C)), DRAFT.

interventions and countermeasures is often measured in terms of lives saved—lives that would be lost to the disease without the intervention.

Vulnerability calculations may also include or aim to capture a number of the costs associated with the contagious disease outbreak—resources required, e.g., medical personnel, hospital beds, and antibiotics; and government interventions, e.g., lost days of work and/or lost taxable revenue during social distancing and the cost of renting and staffing spaces for patient isolation—which in turn can be used to estimate the expected value of economic loss. These costs can be modeled as a function of the numbers of ill and the durations of illness as calculated within the SEIRP model.

There are some costs that should be considered within vulnerability that are far more difficult to capture: (1) trust costs—costs associated with building and/or loss of public confidence in the government’s ability to keep it safe and healthy; (2) social costs—for example, but not limited to, stress- and grief-related costs for adults and children experiencing and witnessing the epidemic; (3) corporate costs—expressed as a function of lost revenue or wages; and others. Even some medical costs are difficult to capture, including costs associated with illnesses and adverse reactions to vaccinations and antibiotics, chronic care for the convalescent, and psychological remedy costs for those not directly impacted. Which losses are considered in the expected value of loss calculations depend on the requirements generated by the decision environment being supported.

There are, therefore, several questions that a decision-maker (e.g., planner or policy maker) focused on return on investment (ROI) might ask when considering how to prepare for and respond to a contagious biological agent epidemic. These questions can be asked for the system as currently established (the status quo) and again with individual and combinations of political, educational, technological, and medical activities initiated to mitigate and/or respond. Each can be evaluated using the proposed vulnerability quantification methodology to estimate costs in terms of lives lost/lives saved and economic impact.

- How many fatalities are expected?
- How many infected persons—who will not become fatalities—are expected?
- How long will medical care be required before victims become fatalities?
 - Determines the amount of time medical services will be required to treat, stabilize, or provide palliative care for those who will become fatalities.
- How long will medical care be required for infected persons who will not become fatalities?
- At what point is the medical system overwhelmed?
 - Contagious biological agent challenges may actually resemble two epidemics: (1) the first epidemic with control measures and available medical care and resources;

and, (2) the second epidemic once the surge capacity for the medical system is overwhelmed and resources become scarce.

- How will the implementation of activities to mitigate the spread of disease before the epidemic begins, change the number of fatalities and infected requiring treatment?
 - Activities may include, among others, public education, vaccination, technological detection and identification deployment, development of additional public health surge capacity, and execution of medicine distribution plans.
- How will the implementation of activities to respond to the spread of disease after the epidemic begins, change the number of fatalities and infected requiring treatment?
- How does delaying these activities, change the number of fatalities and infected requiring treatment?
 - Activities may include, among others, shelter-in-place and mandatory personnel-only policies, community self-distancing, herd vaccination, mass medical distribution, and public health education.

These are just some of the questions that might be asked in the event of a contagious biological event, and, for the most part, they refer only to life safety and incident stabilization. Questions could also be asked about property preservation, the contamination of facilities and resources, the risks to infrastructures drained of resources during an epidemic, and others.

In addition, questions allowing for further differentiation and higher resolution could be posited, although not necessarily answered. Some of these questions include:

- What percentage of the fatalities are elderly? Children? Special needs? Minority, underprivileged, or underserved populations? Socio-economically challenged?
- Subsequently, how could the implementation of targeted policies, procedures, and activities change these percentages?

Although the majority of the questions noted above typically can be answered quantitatively, they may require qualitative assessments as well. For example, when assessing the utility of smallpox vaccination orders, the language of the message and public trust both play a role in determining how much of the population will respond as requested and which communities may respond differently. The data may not exist to quantitatively state whether or not a particular community will respond as directed; the only information available may be anecdotal, however, it is still important to identify and understand its implications.

C. Applying the Guidelines to a Biological Agent Scenario

This appendix proposes the five guidelines listed in section A.2 to assess risk. These guidelines may be applied to any threat and any scenario to provide a quantitative risk assessment for comparison within a single sector or threat area or across sectors. The remainder

of this appendix illustrates how these guidelines would be applied to a contagious biological agent event.

1. Guideline 1: Define Scenarios (Combinations of Attack Vectors/Targets) to Specify the Critical Information Necessary to Estimate Scenario Risk and its Key Parameters—Scenario Consequences, Vulnerability, and Threat.

Risk assessments require the establishment of specific scenarios to provide the essential context for estimating risk as a function of threat, vulnerability, and consequences. The scenario is defined in terms of these three spaces.

a. Threat Definition

The first question is: What is the threat? For our purposes, the threat is the initiation of a contagious biological agent event, either through a naturally occurring epidemic or an intentionally-initiated bioterrorism event. Quantifying the threat requires two additional pieces of information: First, what is the threat agent? And second, what is the route or method of exposure?

The threat agent and likely routes of exposure may be informed by intelligence or may be assumed for planning and risk assessment purposes. Intelligence may come from the intelligence infrastructure or law enforcement sources or, in the case of naturally occurring events, may be gathered through the medical community, historical outbreak information, or medical surveillance and information websites.¹³ The intelligence and medical community may also be able to provide additional information on the likely agent-specific parameters and weaponization/route of exposure parameters; if they are unavailable through intelligence sources, established model parameters may be used or values may be assumed based on available information.¹⁴

What is the threat agent? This is the contagious biological agent that will result in an epidemic; the agent may be naturally occurring—e.g., influenza, SARS—or man-made/weaponized—e.g., plague, smallpox. Once the threat agent is known, the agent-specific parameters may be known or defined:

- Agent infectivity is the likelihood that a specific agent dose and specific exposure route will cause an individual to manifest symptoms;

¹³ The medical community has several resources for conducting surveillance and predicting possible outbreaks and coming epidemics. For example, the *Centers for Disease Control Mortality and Morbidity Weekly Report* provides information on ongoing outbreaks, recent events, and collected health surveillance information. Similarly, ProMed publishes daily and weekly reports on human and animal disease events, as well as weekly and monthly tracking of specific epidemics.

¹⁴ For example, Disraelly, Walsh, and Curling, “A New Methodology,” 2011 provides agent-specific parameters for both plague and smallpox contagious biological casualty estimation modeling.

- Mortality is the likelihood of an agent exposure (or specific dose) to cause a fatality;
- The incubation period is the duration between exposure and manifestation of symptoms;
- The duration of illness/time to death is the time between manifestation of symptoms and convalescence or death during which the individual progresses through the illness stages defined by the injury profile;
- The injury profile defines the illness stages and the associated injury severity levels for both survivors and non-survivors of a contagious biological agent;
- Transmissivity is the probability that the infectious individual can spread the disease to other individuals.

What is the route or method of exposure? For example, in its simplest terms, contagious biological agent events can be initiated through an aerosol release of weaponized agent; through an ingestion contamination incident, usually in food or water; or through the introduction of an index case or infected vector animal into a susceptible population. Of these three possible exposure routes, only the first is likely in an intentional bioterror event; the other two exposure routes may occur naturally or be man-made. As with the agent-specific parameters, values must be defined by the specific type and method of exposure. The following are a few examples of parameters that might be defined for different exposure routes and attack types:

- Aerosol release of weaponized agent through use of a sprayer would include but not be limited to:
 - Sprayer size
 - Rate of release
 - Height of release
 - Agent decay rate
 - Agent particle size
 - *Models for aerosol biological agent release will likely have defined default values for these types of attacks, eliminating the need for the user to define specific parameters.*
- Ingestion contamination of a food source would include, but not be limited to:
 - Food source
 - Quantity/level of contamination (e.g., area of contamination, likely doses associated with normal servings, regions of distribution served by the food producer)
 - Possible routes of agent reduction (e.g., cooking or washing may reduce the quantity of agent ingested)
 - Agent decay rates

- Introduction of an index/multiple index cases into a susceptible population would include but not be limited to:
 - Number of index cases introduced into the population
 - Stage of disease at time of index case introduction—(this determines the level of infectivity)
 - Location of introduction (this can, potentially, influence the transmissivity of an agent; for example, introduction of an index case on an airplane might or might not increase the potential for transmission to a large number of people in a contained space for a long period breathing recirculated air.)
- Introduction of an agent-bearing vector animal into a susceptible population would include but not be limited to:
 - Probability of a vector animal transmitting the agent to a human host
 - Infectivity and transmissivity of the agent among animal vector population
 - Transmissivity of the agent from the animal vector to humans (e.g., how many humans can a single animal vector make ill?)

b. Vulnerability and Consequence Definitions

It is important to know that with a contagious biological event, the vulnerability and the resultant consequences are inextricably linked. Therefore, the questions and parameters necessary to define the scenario in terms of both vulnerability and consequence will be defined together in this section.

The first vulnerability question is: What is the target? As with threat, the answer to the first question dictates the follow-on questions. In addition, the scope of the target considered and the information available and necessary may vary based on who is asking the question and conducting the risk assessment. For example, a risk assessment at the local level may focus on a specific physical location—e.g., a particular part of town that is most likely, like a high-traffic transportation hub or a tourist attraction—or may focus on a specific event—e.g., Fourth of July fireworks and outdoor concert, a popular indoor or outdoor sporting event, a parade. A national risk assessment may ask questions about specific locations or events or may only use these as examples of potential events impacting large swaths of the population and affecting the highest potential for agent and eventual illness dispersion.

Defining the target and its vulnerability also requires some assumptions about the intended consequence of the event: lives lost, property damage, or economic disruption. The likely intended consequence of a contagious biological event is widespread illness resulting in fatalities and casualties; property damage and economic disruption may be secondary results, but they are

not likely the primary intent. Understanding the intent, however, helps determine the parameters necessary to define the target:

- Target population
 - Every day or special event (this may change the number of people in the vicinity of the attack and/or change the distance traveled to attend the event and how far they will travel before the symptoms begin to manifest; this provides information on the potential for multiple simultaneous, seemingly unrelated, outbreaks.)
 - Number of people likely to be in the vicinity
 - Efficacy and distribution of pre-exposure prophylaxis among the population (for example, some years the entire population is offered some form of flu vaccine, whereas other years, the distribution is far more limited.)
 - Distribution of the population in and around the vicinity of the event at the time of initiation
- Physical location
 - Time of day (this may change the rate of agent decay and the transport and dispersion of the agent through an environment.)
 - Meteorology (similarly, depending on the model utilized, this may change the rate of agent transport and dispersion through an environment.)
 - Inside or outside (this provides information on both the potential area through which the agent might be dispersed and, for an event focused even peripherally on physical damage, will inform decisions on denial of space and potential economic loss and restoration costs.)
 - Layers of defense, if any, that might impede the attack and the probability of breaching these layers (it is likely, for a contagious biological agent, that no physical layers of defense exist to prevent or impede an attack or naturally occurring event.)
 - Availability of environmental detection, individual surveillance systems, or other countermeasures (detection, individual surveillance systems, and building protection systems have associated agent levels to trigger detection of a contaminant agent or possible illness and to identify the agent; these levels may then trigger additional actions—e.g., shutting down a building’s ventilation system—or may merely inform public health decisions downstream).

Actions taken during the outbreak, in addition to the existence and implementation of medical prophylaxis, environmental detection, and other protection systems, have the ability to change the overall number of fatalities and casualties resulting from a contagious biological

event. These actions, often referred to as interventions, may be classified as medical and policy-driven or social.

Initiating any intervention depends on identifying the requirement for an intervention, so the first consequence-related question is: How quickly is the outbreak/epidemic/pandemic correctly identified? And, how long does it take for medical authorities to be informed? For example, will interventions be initiated based on an environmental detection and identification, a confirmation, or following the diagnosis of the index case? This decision should be informed by existing policies and procedures.

Medical intervention considerations include the initiation of medical countermeasures, the availability of medical resources, and the efficacy of medical countermeasures initiated before symptom onset and treatment initiated following symptom onset. For each medical intervention considered, the following parameters should be defined:

- Efficacy of the countermeasure or treatment (For what percentage of the population is the countermeasure or treatment effective? What time frame is required? Most vaccines or pre-symptom onset antibiotic prophylaxis courses must be initiated prior to exposure or within a certain time frame after exposure for the countermeasure to be efficacious.)
- When the countermeasure is initiated (When is the countermeasure made available to the population? How long does it take to distribute? This is likely informed by Department of Health and Human Services policy or the policy of a local health department.)
- When the treatment is initiated
- Quantity of the available resource (Whether it is vaccine, antibiotics, hospital beds, ventilators, nurses, or respiratory specialists, the resource is likely to be limited; knowing how much is available may help inform the other interventions necessary to limit the risk of running out of a limited resource or may inform the ways that resources are used. For example, in the event of limited vaccine supplies, ring vaccination—vaccination of the closest contacts of an infectious case and their circle of close contacts—may be selected rather than whole-population vaccination.)
- Percentage of the population receiving and utilizing the intervention (This reduces the susceptible population by an associated fraction. It is important to note that not everyone in the population will comply with any directed intervention and that even among those who comply, there will be some fraction that will not complete the intervention for the required time period—for example, some individuals will pick up pre-symptom onset antibiotics but will not complete the whole course of medication.)
- Percentage of the population suffering adverse effects as a result of medical intervention

Policy-driven or social interventions include social distancing, quarantine of likely exposed and infected, and isolation of infectious patients. For each policy-driven intervention, there are

several parameters that will impact the resulting consequence. Three factors can be defined to allow for ease of evaluation and are likely available or assumed from existing policies and procedures:

- Compliance rate (How many people or what percentage of the population comply with the policy-driven intervention?)
- When intervention is initiated and how long it is maintained (The quicker an intervention is put into place, the more effective it is likely to be. It is worth noting that, as might be expected, the longer an intervention is maintained in place the more effective it is, but, the longer it is maintained, the lower the compliance rate is likely to be.)
- Efficacy of the intervention (For example, social distancing may reduce the number of close contacts, and, therefore, the number of chances for disease transmission that each individual has in a day by an identifiable or an assumed factor.)

c. Scenario Quantification

Once the specific scenario parameters are defined, the contagious biological agent scenario can be quantified. This is usually accomplished by using a series of models and tools. In general, contagious biological scenario quantification involves several steps:

1. A transport and dispersion model is used to determine the spread of the weaponized, aerosolized threat agent as a cloud released from a specific weapon type;
2. The cloud is then applied over a population to determine the dose to each individual or each location within the population;
3. The number of initially exposed and infected individuals is estimated as a function of casualties;
4. A casualty estimation model or tool is then used to estimate casualties and fatalities for the initially exposed population; and
5. The same tool or an additional model may be used to estimate contagious disease spread and the resulting casualties and fatalities.

There are a number of tools to conduct the necessary steps. For example, Vapor, Liquid, and Solid Tracking (VLSTRACK) may be used to model the contagious agent release and cloud dispersion and transport.¹⁵ The Biostrike or Chemical/Biological Strike (CBStrike) models,

¹⁵ Timothy J. Bauer and Matthew G. Wolski, *Software User's Manual for the Chemical/Biological Agent Vapor, Liquid, and Solid Tracking (VLSTRACK) Computer Model, Version 3.1*, (Dahlgren, VA: Naval Surface Warfare Center, Dahlgren Division, 2001).

developed at IDA,¹⁶ could be used to translate the cloud into doses received by a distributed population. The SEIRP methodology can then be used to estimate casualties and disease spread as a function of the initially exposed and infected population. Or, a single package, for example the Hazard Prediction and Assessment Capability (HPAC) software package, can be used to model the contagious biological agent plume, determine the exposures, and estimate population effects.¹⁷ There are also tools to predict the spread of aerosolized contaminants inside buildings and models to predict contamination spread or vector animal transmission.

It should be noted that it is not uncommon for planners to assume an initial number of exposed and infected or a distribution of doses across a population without the use of a specific modeling tools; this option is also acceptable provided that it meets the planners' requirements for rigor and that the assumptions are documented and, as possible, explained and justified.

For the purposes of this illustration, the IDA study team selected the HRIP methodology for casualty estimation. The team is extremely familiar with this methodology and the underlying SEIRP model for contagious biological agent human response. The HRIP contagious biological agent methodology directly implements the quantification of vulnerability by calculating fatalities and casualties over time. In addition, the methodology is versatile in that it allows the user to initiate and incorporate a variety of different interventions, both medical and policy-driven, as well as model contagious disease spread with scientifically derived, peer-reviewed agent parameters.

2. Guideline 2: Ensure that Needs and Limitations of the Supported Decision-making Environment Are Understood and Addressed.

For the purposes of the risk assessment, vulnerability, or *conditional risk*, is quantified as the expected value of loss or consequences given that the scenario occurs. As previously discussed, vulnerability can be quantified as a function of lives lost or as an economic cost. While there are, undoubtedly, economic costs associated with a contagious biological event, the predominant question in quantifying vulnerability for a contagious biological event is: How many lives will be lost? In contrast, the value of most intervention options is measured as a function of how many lives can be saved.

If required, additional calculations can be made to quantify the contagious biological event in economic terms. Fatalities can be correlated to an economic cost using a variety of different estimators. For example, the human capital approach can be used to estimate the value of an

¹⁶ BioStrike and CBStrike are IDA-internally developed models. Additional information on these tools is available from the appendix authors.

¹⁷ Titan Corporation & Defense Threat Reduction Information Analysis Center, *Hazard Prediction Assessment Capability (HPAC) Getting Started*, (Virginia: Defense Threat Reduction Agency, 2004).

individual's life as a measure of future production potential;¹⁸ others have calculated cost of life as a function of both human capital and the costs of wages lost during the duration of illness.¹⁹

Additional costs to consider—and for which actual or estimated dollar values would be needed—include the cost of treatment, the cost of intervention implementation, and the cost of facility recovery and restoration. The cost of treatment can likely be estimated based on historical cases or as a function of the treatment protocols. Treating less lethal forms of influenza with antivirals and bed-rest would be expected to cost less than treating respiratory assistance-requiring plague cases; however, models for both likely exist within the medical community and in the Department of Labor, which retains information on the cost of varying degrees of illness. The cost of intervention implementation is a function of the cost of the intervention, the cost of the people (often volunteers) to distribute the intervention and support intervention efforts, the cost of lost wages, and the cost of lost wages as a result of adverse effects. Finally, the cost of facility restoration and recovery for indoor facilities can potentially be derived from the cleanup efforts following the 2001 anthrax events; a similar procedure would likely be implemented for other biological events. Cleanup, restoration, and recovery of an outdoor space, however, might be harder to estimate.

3. Guideline 3: Use Ratio-scale Metrics to Quantify Threat, Vulnerability, and Consequences.

There are several quantitative conditional risk metrics; the most common include human lives, economic costs, mission impacts, and psychological effects. Other factors can be considered as well, including environmental impacts and loss of public trust, among others. Most commonly, though, as discussed above, risk metrics are expressed in lives and dollars lost.

To ensure that scenarios can be compared, the metrics need to be expressed in ratio scales versus ordinal scales. While ordinal scales allow for a qualitative risk assessment, which may be sufficient to support many decisions, they do not allow for clear, defensible, numeric comparisons across scenarios, options, or sectors.

Following from the previous guidelines, two ratio-scale metrics can be used to assess risk associated with a contagious biological event: lives lost and economic cost. The preferred method, for IDA purposes, is evaluation of lives lost (and possibly additional casualties). Lives lost are clear numerical representations and should be expressed with appropriate estimates of the uncertainty contained in the estimations; overly precise estimates will likely convey false accuracy. The one disadvantage of using only lives lost, however, is that it does not allow for the

¹⁸ J. Steven Landefeld and Eugene P. Seskin, "The Economic Value of Life: Linking Theory to Practice," *American Journal of Public Health* 72, no.6 (1982): 555–566.

¹⁹ B. Cooper and D.P. Rice, "The Economic Cost of Illness Revisited," *Social Security Bulletin* 39 (1976): 21–36.

inclusion of other losses and costs associated with the event, including hospitalization costs, costs of interventions, lost wages and economic revenue, and property restoration and recovery.

The alternate option is to express risk associated with a contagious biological agent event as a function of costs. This requires, as described in the discussion of Guideline 2, the assignment or assumption of costs associated with lost lives and casualties. Additionally, it requires the quantification of costs associated with long-term hospitalizations, intervention implementation, lost wages and economic revenue, and property restoration and recovery. While dollar amounts are a clear numerical representation, they should still be represented in thousands or even millions of dollars to preclude false accuracy.

Ratio-scaled estimates could also be used for comparison purposes; these expected values could also be expressed as ratios of lives lost and/or cost of consequences post-intervention over the baseline estimates.

4. Guideline 4: Quantify Vulnerability Estimates as an Expected Value of Loss.

Conditional risk is the expected value of loss or consequences. In an event where a layered defense exists and there is a distinct before-event period, for which vulnerability can be estimated, and after an event period, for which consequences can be calculated, conditional risk is the joint probability of successfully penetrating all defensive layers multiplied by the consequences.

As discussed in the main portion of this document, conditional risk is an appropriate representation of vulnerability when the computation of joint probability is too complex and intractable to be calculated by other methods. As noted above, the vulnerability and consequences of a contagious biological event are linked and cannot be easily defined independently, in part because the exact moment of the event is not easily defined. The moment of the attack could be defined as the moment at which the attack occurs, the moment the first patient manifest symptoms, the point at which the disease is positively diagnosed, the moment at which a threshold number of patients is surpassed or any other potential point in the course of disease.

Because a contagious biological event is ongoing and continuously evolving, there is no clear before—vulnerability—and after—consequence. As long as there remains some portion of the population that is infectious and transmitting the disease, some portion of the population remains susceptible. Additionally, interventions that will help mitigate the effects in those already exposed and infected will help prevent the transmission of disease to those who are susceptible. Thus, for a contagious biological agent event, as with many cascading events, there is no clear delineation between pre-event vulnerability and post-event consequence for calculation purposes.

As described in Guideline 2, the HRIP contagious biological agent methodology is used in this appendix to estimate fatalities and casualties—the expected value of loss—for specific

scenarios without interventions and then with interventions. This allows for a clear comparison of the decreased conditional risk associated with individual (or combinations of) interventions expressed as decreased lives lost (lives saved) calculated by subtracting the number of fatalities with a specific intervention (i.e., Intervention 1) from the number of fatalities with no interventions as shown in equations B-1–B-3.

$$\begin{aligned}
 \text{ConditionalRisk}_{\text{NoIntervention}} &= \text{LivesLost}_{\text{NoIntervention}} \\
 \text{ConditionalRisk}_{\text{Intervention1}} &= \text{LivesLost}_{\text{Intervention1}} \\
 \text{Decrease in ConditionalRisk}_{\text{Intervention1}} \\
 &= \text{LivesLost}_{\text{NoIntervention}} - \text{LivesLost}_{\text{Intervention1}} = \text{LivesSaved}_{\text{Intervention1}} . \quad \text{eqs. B-1–B-3}
 \end{aligned}$$

Similar calculations could be conducted as a function of economic loss and costs associated with a contagious biological agent event.

It should be noted that, depending on the needs of the decision-maker or planner, conditional risk can be expressed in several ways, including *worst case*, *probability weighted*, *best case*, and *bounded range*. Any expression may be appropriate as long as it is clearly defined by the planner along with the appropriate scenario and parameters, then used consistently; for example, it would not be appropriate to compare the conditional risk, or expected value of lives lost, and no intervention in a best case scenario with the conditional risk associated with the use of post-exposure vaccination. Because the two scenarios are not commensurate—one being worst case and the other being best—they may not be comparable. A comparison could still be made, because both sets of results are calculated in terms of the expected value of lives lost, but any comparison of these results should clearly indicate both the scenario and the type of conditional risk being considered.

5. Guideline 5: Manage the Scenario Space so that the Scope, Scale, and Assumptions can be Verified as Appropriate to the Decisions, or Modified, if Determined to be Indefensible.

To conduct an effective risk assessment and make a decision appropriate to the decision context being considered, the planner needs to manage the scenario space and choose scenarios appropriate to the risk assessment. Managing the scenario space requires a return, in part, to Guideline 1:

- What is the perceived threat?
- What is the perceived target?

If the answer to these questions is very specific, then the scenario space is easy to manage. For example, if the concern is the release of a specific contagious biological agent by an individual wielding a sprayer in Times Square on New Year’s Eve, then the perceived threat is a man-portable sprayer filled with the specific agent, and the perceived target is the population in

Times Square. A man-portable sprayer that would not be immediately obvious can be only so big (e.g., small enough to carry in a shopping bag or backpack) and so the volume is likely limited to a certain, assumable range, and if the agent is known, then certain weaponization parameters may be assumable as well. There are several places a release could happen in Times Square, but three or four model runs (and possibly fewer) would likely provide a representative sample of the likely human response and estimated casualties resulting from the event.

Conversely, if even one variable in that scenario changes, then the number of permutations and the scope of the scenario space increase significantly. For example, if the perceived threat is any contagious biological agent (versus a specific agent), then the modeler needs to determine how many and which agents to consider. Now, instead of three runs, it may be three runs times the three most likely weaponized biological warfare agents.

The scenario space can be as immense as the scenarios the decision-maker and planner can imagine; therefore, the decision-maker and planner need to ensure that the scenario space is informed by external information, where available, and, ideally, representative of the possible scope of threat events.

The more vague the questions, therefore, the more information that needs to be provided to define the scenario space:

- What assumptions govern the scenario space?
 - The parameters that were defined in Guideline 1 are assumptions and so are already defined.
 - Assumptions for clarification must include the assumptions of each associated model.
- What are the limitations of the scenario space?
 - Limitations in scale, scope, and spread must be addressed; for example, a scenario that will not account for spread occurring outside the local area must be documented.
 - The original HRIP methodology did not include treatment; this would be an example of a limitation of the methodological approach.
 - Limitations for clarification must include the limitations of each associated model.
- What is the scope of the scenario space?
 - Scope includes:
 - Breadth of the attack (i.e., how many locations need to be considered in the model—cities, places in each individual city, smaller or more rural locations)
 - Methods of executing the attack

- Attack agent being used (typically a set of comparable scenarios is established for a single agent)
 - Which cases are being modeled (e.g., best case, worst case, range of cases; if the intention is to compare intervention efficacies, then it may be suitable to model only the worst case scenarios.)
 - How wide the attack will be allowed to spread (while contagious biological agent events have the potential to turn global, most decision-makers are concerned with a more limited geographic area and can scope their scenarios to that region.)
- The scope is likely defined by the planner or decision-maker; a planner in a large east coast city is probably less worried about attacks on the west coast but may still choose to model the potential for intentional or inadvertent human vectors introducing a contagious biological agent following a west coast event. If the decision-maker has already modeled a human vector, or even only an intentional larger scale release, then modeling a west coast attack may not be necessary.
 - As noted above, scope should be informed, as much as possible, by information available from the intelligence community, the Departments of Health and Human Services and Homeland Security, local health departments, and other available sources.

To fully evaluate the implications of a contagious biological agent event, multiple analytical cases will be required including variations and combinations of attack locations, population size, disease spread, interventions, times of intervention initiation, and intervention efficiency. (For example, some populations will refuse to comply with a vaccination; this results in continued possible spread and duration of the disease in some areas.)

6. Applying the Guidelines

Applying the guidelines to define a range of scenarios is an iterative process. In part, this is why establishing the scope early on is important. Each scenario will likely suggest other possible scenarios for modeling. If, however, the field of consideration—possible locations, likely attacks, target types—is identified early on, it should help the decision-maker and the planner limit the number of scenarios and therefore the number of runs.

These guidelines are intended to help guide decision-makers and planners in the development of scenarios and quantitative risk assessments for contagious biological agent events. They are intended to serve as a guide and should not be considered all inclusive; other assumptions, parameters, and models may and should be considered depending on the goals, scope, and scenarios of interest to the decision-maker.

D. Biological Agent Example

Section C identified the framework for applying the risk assessment quantification guidelines to a contagious biological agent event. To help the decision-maker and planner understand how to apply the framework, this chapter will provide a quantitative example of the application of the guidelines to a contagious biological agent scenario.

For the sake of simplicity, the scope of the scenario space will be kept relatively small. Further, to avoid the use of multiple modeling tools, an initial number of exposed and infected will be assumed. For the purposes of this example, the IDA team assumes that a local decision-maker or planner is attempting to quantify the risk associated with an aerosolized, weaponized plague attack on a local bus station. The attack remains relatively local (which although potentially unlikely, has happened in naturally occurring outbreaks; for example, SARS, although it produced individual cases in multiple cities, only resulted in a small number of epidemic centers) but might represent the scenario a local decision-maker is actually concerned about.

For the purposes of this contagious biological agent example, the intentional release of aerosolized contagious biological agent results in exposure to some fraction of a closed population—e.g., a city with a limited, fairly constant, population. The population is generally in good health, and no one within the population to be exposed has been vaccinated against the agent prior to the release.

This provides the starting point for the scenario. To make it easier to follow this example and potentially follow a similar path, parameters and other values are compiled in tables.

1. Guideline 1: Define Scenarios (Combinations of Attack Vectors/Targets) to Specify the Critical Information Necessary to Estimate Scenario Risk and its Key Parameters—Scenario Consequences, Vulnerability, and Threat.

a. Threat Definition

The first question is: What is the threat?

The threat is an intentional bioterrorism event: an aerosolized, weaponized plague release from a man-portable sprayer. This identifies the threat agent and the route of exposure—the agent is aerosolized to cause inhalation exposure.

Once the threat agent is known, the agent-specific parameters may be known or would be defined as required by the selected biological agent contagious spread model. Likely parameters would include infectivity, incubation period, lethality, duration of illness, and the efficacy and duration of prophylaxis.

As with the agent-specific parameters, values must be defined by the specific type and method of exposure. Because, for ease of demonstration, the IDA study team has assumed a

particular attack yield, it does not need to specify weaponization parameters. For reference, however, Table B-1 is provided as a sample. Many of the values in Table B-1 could be drawn from the default values for a transport and dispersion model (i.e., VLSTRACK).

Aerosol release of weaponized agent through use of sprayer parameters would include but not be limited to:

Table B-1. Weaponized Agent Parameters Summary Table

Parameters
Sprayer size
Rate of release
Height of release
Agent decay rate (day)
Agent decay rate (night)
Agent particle size

b. Vulnerability and Consequence Definitions

The first vulnerability question is: What is the target?

The target is a local bus terminal as identified above. Additional vulnerability parameters are defined in Table B-2. The consequences or loss will be defined in terms of lives lost, so parameters will focus on numbers of individuals potentially affected by the threat.

This table is intended for demonstrational purposes only. It is not intended to represent any specific location or city. All values are notional.

Table B-2. Target Population Parameters Summary Table

	Parameters
Target population	Rush hour travelers
Every day/special event	Every day, evening rush hour
Number of people in the vicinity ^a	1500
Pre-exposure prophylaxis—efficacy	0
Pre-exposure prophylaxis—distribution	0
Distribution of the population at t=0	Evenly distributed across the space
Distribution of the population post-event	All personnel remain approximately local to the event
Number of anticipated distributed outbreaks ^a	1 (all personnel remain local)

^a Notional value, provided for illustrative purposes only.

There is no approved pre-exposure prophylaxis for plague, so the efficacy and distribution of the pre-exposure prophylaxis are both zero.

The parameters defining the physical location are listed in Table B-3. Again, because, for the purposes of this example, no external model is being used to estimate the exposures, the table is for informational purposes only and no values are included. It should be noted that because the event is indoors, at a bus station, meteorological conditions and time of day would not have as significant an impact on the transport and dispersion of the agent as it would if the event occurred outside.

Table B-3. Target Physical Location Parameters Summary Table

	Parameters
Indoors or Outdoors	Indoors
Time of day	
Meteorology—Temperature	
Meteorology—Wind speed	
Meteorology—Wind direction	
Meteorology—Relative humidity	
Layers of physical defense	None
Availability of environmental detection	Not available

Interventions initiated during the outbreak can change the overall number of fatalities and casualties resulting from a contagious biological event. Initiation of these interventions in the model may be considered new scenarios or derivative scenarios from the worst-case, no intervention scenario. For the purposes of this example, the IDA study team will consider two interventions—medical post-exposure, pre-symptom onset antibiotic prophylaxis and social-distancing. Sample medical intervention parameters that might be used for modeling these interventions are shown in Table B-4.

Table B-4. Medical Intervention Parameters Summary Table

	Post-Exposure, Pre-Symptom Onset Antibiotic Prophylaxis Parameters
Pre-exposure prophylaxis efficacy	
Post-exposure prophylaxis efficacy	
First day of prophylaxis administration	
Last day of prophylaxis efficacy	
Quantity of prophylaxis available ^a	Unlimited
Prophylaxis Compliance rate ^a	75%
Prophylaxis adverse reaction rate ^a	0%
Treatment efficacy	
Treatment initiation	
Treatment duration	
Quantity of treatment available	
Treatment adverse reaction rate	

^a Notional value, provided for illustrative purposes only.

The policy-driven intervention parameters are shown in Table B-5. Because quarantine and isolation are not considered in this example, the associated parameter values are left blank.

Table B-5. Policy-driven Intervention Parameters Summary Table

	Social Distancing Parameters	Quarantine Parameters	Patient Isolation Parameters
Intervention Compliance rate ^a	100%		
Intervention initiation day post-exposure ^a	7		
Duration of intervention ^a	30 days		
Efficacy of intervention ^a	Reduces contacts by a factor of 4		

^a Notional value, provided for illustrative purposes only.

c. Scenario Quantification

Once the specific scenario parameters are defined, the contagious biological agent scenario can be quantified. This is, usually, achieved through the use of a series of models and tools. Quantifying the general contagious biological scenario involves several steps, which are summarized in Table B-6.

Table B-6. Steps for Contagious Biological Scenario Quantification

Steps	Tools
1. A transport and dispersion model is used to determine the spread of the weaponized, aerosolized threat agent as a cloud released from a specific weapon type	Assumed (e.g., VLSTRACK, HPAC)
2. The cloud is then applied over a population to determine the dose to each individual or each location within the population	Assumed (e.g., CBSTRIKE, HPAC)
3. The number of initially exposed and infected individuals is estimated as a function of casualties	Assumed (e.g., HRIP, HPAC)
4. A casualty estimation model or tool is then used to estimate casualties and fatalities for the initially exposed population	HRIP
5. The same tool or an additional model may be used to estimate contagious disease spread and the resulting casualties and fatalities	HRIP

It should be noted that it is not uncommon for planners to assume an initial number of exposed and infected, which IDA has indicated by inserting the word “Assumed” above.

For the purposes of this evaluation, IDA selected the HRIP methodology for casualty estimation. The methodology is versatile in that it allows the user to initiate and incorporate a variety of different interventions, both medical and policy-driven, as well as model contagious disease spread with scientifically derived, peer-reviewed agent parameters.

2. Guideline 2: Ensure that the Needs and Limitations of the Supported Decision-making Environment Are Understood and Addressed.

For the purposes of the risk assessment, vulnerability or conditional risk, is quantified as the expected value of loss or consequences given that the scenario occurs. For the purposes of the example, vulnerability will be quantified as a function of lives lost.

The HRIP contagious biological agent methodology quantifies vulnerability by calculating fatalities and casualties over time.

3. Guideline 3: Use Ratio-scale Metrics to Quantify Threat, Vulnerability, and Consequences.

The ratio-scale metrics for lives lost would be a numerical representation (likely expressed in tens or hundreds to minimize false accuracy), as shown in Table B-7.

4. Guideline 4: Quantify Vulnerability Estimates as an Expected Value of Loss.

Conditional risk is the expected value of loss (or consequences); as noted in the discussion of Guideline 2, conditional risk is defined for this scenario as the expected number of lives lost or fatalities.

To effectively calculate the baseline expected value of loss, there are additional parameters that need to be defined for use in the HRIP methodology. These parameters are shown in Table B-7; it should be noted that all parameters are notional and are not intended to reflect any specific city or scenario.

Table B-7. Initial Conditions Summary Table

	Initial Conditions
Population at Risk^a	50000
Population for whom pre-exposure prophylaxis is efficacious^a	0
Susceptible population	49000
Exposed and infected population^a	1000
Infectious Stage 1 population^a	0
Infectious Stage 2 population^a	0
Total removed population	0
Fatalities	0

^a Notional value, provided for illustrative purposes only.

The notional baseline expected value of loss is the number of lives lost as a result of the initial scenario. For this example, approximately half the population, or 25,000 people, become fatalities when no medical intervention or treatment is provided, as shown in Table B-8.

**Table B-8. Baseline Expected Value of Loss for Notional Scenario
(Number of People Affected)**

	7 days	15 days	30 days	60 days
Prompt fatalities	0	0	0	0
Delayed fatalities	700	5700	23200	25000
Casualties	900	3200	700	0
Convalescence	0	0	0	0
Asymptomatic	48400	41000	26100	25000

When post-exposure, pre-symptom, onset antibiotic prophylaxis is initiated on Day 7 for seven days, without medical treatment, the expected value of loss drops to approximately 6,800 people, as shown in Table B-9.

Table B-9. Baseline Expected Value of Loss for Notional Scenario with Medical Antibiotic Prophylaxis at Day 7 for 7 Days (Number of People Affected)

	7 days	15 days	30 days	60 days
Prompt fatalities	0	0	0	0
Delayed fatalities	700	5000	6600	6800
Casualties	900	900	100	0
Convalescence	0	0	0	0
Asymptomatic	48400	44100	43300	43200

Following the initiation of social distancing on Day 7 without medical treatment, the expected value of loss is 7,500 fatalities as a result of the exposure, as shown in Table B-10.²⁰

Table B-10. Baseline Expected Value of Loss for Notional Scenario with Social Distancing at Day 7 (Number of People Affected)

	7 days	15 days	30 days	60 days
Prompt fatalities	0	0	0	0
Delayed fatalities	700	4400	7400	7500
Casualties	900	1000	100	0
Convalescence	0	0	0	0
Asymptomatic	48400	44600	42500	42500

The initiation of additional interventions, including medical treatment, would further reduce the expected value of loss.

5. Guideline 5: Manage the Scenario Space so that the Scope, Scale, and Assumptions can be Verified as Appropriate to the Decisions, or Modified, if Determined to be Indefensible.

Typically, numerous runs would be required to fully evaluate the implications of this type of incident—including variations and combinations of attack locations, population size, disease spread, interventions, times of intervention initiation, and intervention efficiency. Because the scope of this particular scenario was so narrow, however, only the limited number of runs that appear in Table B-10 are necessary.

To facilitate managing the scenario space, the IDA study team addressed the following questions associated with Guideline 5 below.

²⁰ For more information on implementing medical and policy-driven interventions in the HRIP contagious biological agent methodology, please contact Deena S. Disraelly.

- What assumptions govern the scenario space?
 - The parameters, as identified in Guideline 1, are the assumptions that dictate the agent type, the threat, the target, and the interventions.
 - Additionally, the parameters identified in Guideline 3 are the assumed, notional initial conditions for this scenario.
 - The assumptions associated with the HRIP contagious biological agent methodology are covered in detail in Disraelly, Walsh, and Curling, “A New Methodology.”
 - Assumptions for clarification must include the assumptions of each associated model.
- What are the limitations that govern the scenario space?
 - For the modeled, notional scenarios, no medical treatment was considered.
 - The modeled, notional scenarios do not allow for expansion or transmission of the contagious biological agent beyond the regional boundaries of the scenario.
- What is the scope of the scenario space?
 - The scope includes:
 - Breadth of the attack—single location, notional local bus station
 - Methods of executing the attack—single method, man-portable sprayer filled with specific contagious biological agent
 - Attack agent being used—Plague
 - Which cases are being modeled—single notional case with two interventions—medical, post-exposure, pre-symptom onset antibiotic prophylaxis; and policy-driven social distancing
 - How wide the attack will be allowed to spread—for the purposes of this scenario, the attack is not allowed to spread beyond the single, initial, isolated city.

E. Conclusions

The objective of this appendix is to illustrate how the definitions and guidelines described above can be implemented to assess and estimate vulnerability and risk in a biological threat context—both naturally occurring and from man-made biological agents. The intent is to show how estimates for both risk and vulnerability as a result of a complex biological event, can be formed in such a way that the results are commensurable with risk and vulnerability estimates computed in other critical infrastructure sectors using the definitions and guidelines above. The

IDA study team noted, however, that the methods of estimation can vary significantly across sectors and yet still achieve the same desired outcome of commensurability. Biological events, like many other events, are potentially cascading, involving risk to systems and sectors beyond the general population that becomes infected as a result of the disease.

It is important to note that this appendix provides a great deal of information, yet barely scratches the surface of what is, potentially, an extremely challenging subject for risk assessment. Whereas many hazardous events, e.g., terrorist events and natural disasters, occur at a distinct moment in time, contagious biological agent events are continuous and constantly evolving.

And, while most hazardous events have their direct effects concentrated in a single physical location, even one as large as a state or region, a contagious biological agent has the potential to produce casualties and impacts far beyond the initial event area, even before the first symptoms are identified and diagnosed.

Additionally, whereas many disasters have physical layers of protection to mitigate effects, defensive measures for contagious biological events are often initiated after initial exposure occurs but before everyone is exposed and infected. As a result, the delineation between vulnerability to an event and consequences after an event is blurred.

Further, decision-makers and planners can prepare for a contagious biological event locally; however, the reach of these diseases, as well as the likelihood of competition for scarce resources, makes risk assessments conducted in a local vacuum less likely to be realistic or effective.

That is not to say, however, that decision-makers or planners should not plan for these events. To the contrary, decision-makers and planners can use tools like the risk assessment framework proposed here and the HRIP contagious biological agent methodology to conduct assessments and comparisons of potential interventions—especially policy-driven and social interventions that limit the need for resources but have high potential for contributing to the mitigation and cessation of a contagious biological agent event.

Appendix C

Estimating Vulnerability in the Information Technology Supply Chain

Published separately

Appendix D

Ordinal Scales and Risk Assessment

Although multiple methods can be used to estimate risk, they can be reduced to two categories: (1) qualitative and (2) quantitative. Both categories require measurement scales to estimate values for the key risk variables—threat, vulnerability, and consequences. This appendix discusses the use of ordinal and ratio scales¹ to estimate risk parameters, and methods of combining threat, vulnerability, and consequence estimates using these scales.

Given that the objective of its task was to provide *quantitative* estimates of vulnerability, the Institute for Defense Analyses (IDA) study team found that certain methods and measurement scales were necessary to achieve this objective while others either failed or were unsupportable. Specifically, ordinal scales pose well-known problems of misinterpretation and poor resolution, and are often misused to produce risk metrics (e.g., in the form of *risk matrices*) that are unsupportable or misleading. Moreover, using ordinal scales to measure key risk variables can compound uncertainty when they are used to estimate risk as an expected value. Even when used properly in risk assessments, ordinal scales provide only limited information to decision-makers.

The material in this appendix draws heavily from two sources that the reader is encouraged to consult for additional discussion and examples:

- Louis Anthony Cox, Jr., “What’s Wrong with Risk Matrices?” *Risk Analysis* 28, no. 2 (2008): 497–512.
- D. Hubbard and D. Evans, “Problems with scoring methods and ordinal scales in risk assessment,” *IBM Journal of Research and Development* 54, no. 3, Paper 2, May/June 2010.

¹ An *ordinal scale* is a measurement scale that assigns numbers (or other labels) to attributes so that the relative order of the labels is meaningful but differences or ratios are not meaningful. In contrast, a *ratio scale* assigns numbers to attributes so that differences and ratios are meaningful and there is a natural zero point to the scale. Further background on measurement scales is provided later in this appendix.

A. Background

1. Stevens' Scales of Measurement

In 1946, S.S. Stevens introduced a hierarchy of four types of measurement scales: nominal, ordinal, interval, and ratio.² A measurement scale assigns numbers (or in the case of nominal and ordinal scales, perhaps other labels) to attributes. The different types of scales arise from different meaningful operations. Briefly, in a nominal scale, only equality is meaningful; in an ordinal scale, relative order (that is, a greater than/less than comparison) is also meaningful; in an interval scale, taking differences between assigned numbers is also meaningful; and finally, in a ratio scale, all of the previous operations, as well as forming ratios between assigned numbers, are meaningful.

Table D-1 lists some examples of common measurement scales, including nominal, ordinal, ratio, and interval. This appendix focuses on ordinal scales (and, to a lesser extent, ratio scales), due to their prevalence in risk assessments.

Table D-1. Common Examples of Measurement Scales

Nominal	Ordinal	Interval	Ratio
Gender	Movie ratings (0, 1, or 2 thumbs up)	Degrees F	Degrees K
Ethnicity	USDA quality of beef (good, choice, prime)		Annual income in dollars
Categorization of rocks as igneous, sedimentary, or metamorphic	Hardness scale of minerals		Length or distance in centimeters, feet, miles, etc.
	The rank order of anything		

Note: USDA = United States Department of Agriculture

Stevens also defined the following characteristics for each type of measurement scale:³

- Basic empirical operations
- Mathematical transformations that leave the scale invariant
- Permissible statistics on data acquired from the scale

² S.S. Stevens, "On the theory of scales of measurement," *Science* 103 (1946): 677–680. Also see S.S. Stevens, "Mathematics, Measurement, and Psychophysics," in *Handbook of Experimental Psychology*, ed. S. S. Stevens (New York: John Wiley & Sons, 1951).

³ Stevens, "On the theory of scales of measurement."

From the basic empirical operations one can derive which arithmetic operations are meaningful for a particular type of scale. These characteristics are summarized in Table D-2. The empirical and arithmetic operations and permissible statistics are cumulative: the operations and permissible statistics appearing in each row also apply to scales listed in succeeding rows. For instance, ordinal scales have operations for determining equality and rank order, and permissible statistics for ratio scales include all of the statistics listed in the last column.

Table D-2. Summary Features of Measurement Scales⁴

Scale	Empirical Operations	Permissible Arithmetic Operations	Transformations Which Leave Scale Invariant	Permissible Statistics
Nominal	Determination of Equality	Counting	One-to-one functions	Number of Cases, Mode
Ordinal	Determination of greater or less than	Greater than/less than	Monotonic increasing functions	Median, Percentiles
Interval	Determination of equality of intervals or differences	Addition, Subtraction	Multiplication by a positive constant and/or addition of a constant	Mean, Standard deviation
Ratio	Determination of equality of ratios	Multiplication, Division	Multiplication by a positive constant	Coefficient of variation

The strongest permissible arithmetic operations for variables measured using ordinal scales are *greater than—less than* comparisons. It is possible to say that the value of one variable is greater than another, but it is not possible to say *how* much greater. For instance, a rating of 4 on an ordinal scale measuring pain is *not* considered twice as painful as a 2 rating on the same ordinal scale, but the rating of 4 is greater than the rating of 2. Consequently, differences between levels in an ordinal scale are not meaningful, and operations that require computing differences, distances, or ratios of values also lack meaning when those values are measured using ordinal scales. The only valid operations are ones that determine equality or a relative comparison of magnitudes (greater than or less than). Since ordinal scales are rank ordered, they are invariant under transformations by monotonic increasing functions (since such functions preserve the rank order).⁵ Stevens asserted that ordinal scales provide only a handful of permissible statistical computations, including the number of cases, mode, median, and percentiles. With regard to statistical operations applied to variables measured using ordinal scales, he contended that “in the strictest propriety the ordinary statistics involving means and standard deviations ought not to be used with these scales, for these statistics imply a knowledge of something more than the relative rank-order of data.”⁶

⁴ Adapted from Stevens, “On the theory of scales of measurement.”

⁵ A function f is *monotonic increasing* if $f(a) > f(b)$ whenever $a > b$.

⁶ Stevens, “On the theory of scales of measurement.”

Adding, subtracting, multiplying, or dividing two variables measured using ordinal scales are not considered permissible. Hence, when combining threat, vulnerability, and consequence variables measured using ordinal scales to compute a risk metric, it is improper to combine them by using a product, sum, difference, or quotient. Accordingly, forming risk as a product of threat, vulnerability, and consequences is not permissible when using variables that are estimated with ordinal scales.

Ordinal scales vary significantly in resolution and construct, and two specific types are often applied in risk analysis models. This appendix defines a *pure ordinal scale* as a set of discrete, qualitative labels in which order is implied. For instance, a Likert scale⁷ that includes the phrases {Strongly Agree, Agree, Neutral, Disagree, Strongly Disagree} is a pure ordinal scale, because an order is implicit in the language. A *threshold ordinal scale* is constructed from a ratio scale that has been subdivided into *bins* with a rating (qualitative or quantitative) given to each bin. As an example, consider the variable Probability of Winning, which, by the axioms of probability, must fall in the interval [0, 1]. A threshold ordinal scale might consider Low to be [0, 0.3], Medium to be [0.3, 0.7], and High to be [0.7, 1]. For both pure ordinal scales and threshold ordinal scales, Stevens' statements about the lack of meaningfulness of differences hold. In the pure ordinal scale, the difference between Neutral and Agree is undefined. In the threshold ordinal scale, the difference between Low and Medium could be any value between 0 and 0.7 (70% of the entire range of the variable).

2. Discussion of Stevens' Measurement Scale Taxonomy

Stevens' assertions about permissible statistics on scales of measurement generated both proponents and opponents in the fields of measurement theory and data analysis. Many of Stevens' ideas were promulgated in statistics textbooks, and they influenced the development of the field of non-parametric statistics.⁸ The principal assertion continued to be that there are fundamental differences between ordinal scales and other types of scales. Thomas P. Wilson, confirming Stevens' assertions in "Critique of Ordinal Variables," stated that when only ordinal measurements are possible, only very weak conclusions, if any, could be drawn from the data relative to a proposed model.⁹ His argument details the restrictions ordinal variables place on model development, which, when ignored, can produce results "completely dependent on how the variables were originally coded."¹⁰

⁷ Likert scales, named after its inventor, psychologist Rensis Likert, measure attitudes by asking people to respond to a series of statements about a topic, in terms of the extent to which they agree with them (e.g., strongly agree, disagree, etc.)

⁸ S. Siegel, "Nonparametric Statistics," *The American Statistician* 11 (1957): 13-19; H. M. Marcus-Roberts and F. S. Roberts, "Meaningless Statistics," *Journal of Educational Statistics* 12 (1987): 383-394.

⁹ Thomas P. Wilson, "Critique of Ordinal Variables," *Social Forces* 49 (1971): 432-44.

¹⁰ Ibid.

Many of Stevens' critics contend that an ordinal variable is a crude representation of an underlying interval or ratio variable (this is what was defined above as a threshold ordinal variable) and discuss situations where ordinal variables can be treated as interval with minimal error, but only when the model is designed carefully and the analysis is precise.¹¹ Jarl Kampen and Marc Swyngedouw raise several objections to this stance, notably the lack of proof that such an underlying variable exists and the inability to identify its distribution to determine thresholds.¹² Further, if an ordinal scale is simply a representation of a higher-order scale and uncertainty is involved in variable determination, why bother with the ordinal scale at all? Overall, both proponents and opponents of Stevens' work advocate for careful, thoughtful development of models with an acute awareness of the limitations that ordinal variables can create in the output variables.

B. Prevalence of Ordinal Scales in Risk Models

The simplicity of ordinal scales accounts for their prevalence in government organizations. When agencies estimate the risk posed by uncommon, large-scale events, where little or no data are available, a model composed of multiple ordinal scale variables (or risk matrix) provides a structured decision rubric that is easy to understand and accessible to all stakeholders in a project. Color-coding schemas and compact information displays are often incorporated with ordinal scales to make successful briefing and communication tools.¹³ Moreover, many users of ordinal scales avoid the use of more quantitative (e.g., ratio) scales because they are reluctant to quantify estimates of uncertain variables.

However, the use of ordinal scales in such instances does not alleviate problems of uncertainty or limited empirical data; on the contrary, it obscures the uncertainty and adds an additional level of potential errors during variable estimation.¹⁴ The utility (and popularity) of ordinal scales lies in their ability to differentiate high risks from low risks, but their utility and robustness in the middle ranges or for calculating the return on investment of various decision options is limited.¹⁵

Unfortunately, data supporting (or not supporting) the effectiveness of ordinal-scaled models in real decision environments are very limited.¹⁶ Accordingly, in many instances, agencies turn to ordinal scale risk models, not because of their proven effectiveness, but because

¹¹ See, for example, the "underlying variable approach" discussed in Jarl Kampen and Marc Swyngedouw, "The Ordinal Controversy Revisited," *Quality and Quantity* 34 (2000):87–102.

¹² Ibid.

¹³ Louis Anthony Cox, Jr., "What's Wrong with Risk Matrices?" *Risk Analysis* 28 (2008): 497–512.

¹⁴ Douglas W. Hubbard, *The Failure of Risk Management* (Hoboken, NJ: John Wiley & Sons, 2009).

¹⁵ Cox, "What's Wrong with Risk Matrices?"

¹⁶ D. Hubbard and D. Evans, "Problems with scoring methods and ordinal scales in risk assessment," *IBM Journal of Research and Development*, Vol. 54, no. 3, Paper 2, May/June 2010.

of their simplicity and their apparent applicability across many risk domains. For example, the National Institute of Standards and Technology (NIST) recommends using an ordinal rating system comprised of likelihood (a combination of threat and vulnerability) and impact, each measured using a three-tiered ordinal scale (high, medium, low), in order to assess the risk of information security breaches. The NIST *Risk Management Guide for Information Technology Systems* contains a brief discussion of the advantages and disadvantages of using qualitative ordinal scales:

The main advantage of the qualitative impact analysis is that it prioritizes the risks and identifies areas for immediate improvement in addressing the vulnerabilities. The disadvantage of the qualitative analysis is that it does not provide specific quantifiable measurements of the magnitude of the impacts, therefore making a cost-benefit analysis of any recommended controls difficult.¹⁷

It should be noted that the NIST *Guide* may overstate the ability of a qualitative ordinal rating system to prioritize risks. Simple yet defensible quantitative risk estimates often expose the inconsistencies and arbitrary nature of qualitative risk estimates and prioritizations derived from them.¹⁸

Other organizations profess similar philosophies. The Department of Energy Risk *Management Guide* states:

During the qualitative analysis, the probability and consequence scales can be categorical. However, it is often useful to assign quantitative metrics to the qualitative categories to help ensure consistent assignment of probabilities and consequences across a project/program.¹⁹

Documents addressing risk management throughout the Government—including military services, food and drug protection, and virus control—often reflect preferences for these types of qualitative ratings.

C. Problems with Using Ordinal Scales in Risk Models

1. General Concerns

There is a long contested debate over the reliability and consistency of verbal probability cues such as High, Medium, Low, Likely, Unlikely, Very Unlikely. Substantial variability exists between how subjects interpret the numerical values assigned to these phrases. This variability

¹⁷ Gary Stoneburner, Alice Goguen, and Alexis Feringa, *Risk Management Guide for Information Technology Systems*, NIST 800-30 (Gaithersburg, MD: National Institute of Standards and Technology, July 2002).

¹⁸ Louis Anthony Cox, Jr. “Some Limitations of Qualitative Risk Rating Systems,” *Risk Analysis* 25 (2005): 651–662.

¹⁹ Department of Energy, *Risk Management Guide*, DOE G 413.3-7 (Washington, DC: September. 16, 2008).

persists even when subjects are given explicit, quantitative instructions on how to interpret these labels.²⁰

The discrete values used in ordinal scales often cause a loss of information due to their poor resolution and the related phenomenon of range compression. Range compression occurs when a common, often qualitative, label is applied to values that have significant quantitative differences. When variables are assessed by categorizing them into a small number of bins, significantly different values may be treated as equivalent (e.g., adverse events that occur with probabilities 0.001 and 0.1 may be placed in the same Unlikely bin).²¹ Range compression errors are exacerbated when most of the estimated variables fall into a small number of the possible bins (e.g., when assessing variables by binning them into five categories, subjects may place most of the variables into just two of the bins), and there is evidence that these situations are not uncommon.²² Moreover, if values assessed using a threshold ordinal scale are later translated to quantitative values (e.g., by representing each ordinal range by a single point value) and this quantitative intention is not described in the original assessment, the analyst is inferring data that were not present in the survey of the assessors.

There are additional errors that have the potential to affect all risk assessments (regardless of how quantitative they are and the measurement scales they use). These may be exacerbated or more difficult to correct when ordinal scales are used. Among these are various cognitive biases (such as confirmation biases, optimism biases, and overconfidence). While quantitative methods can be calibrated to correct for these cognitive biases, qualitative risk assessments that use ordinal scales typically include no such calibration to account for these known biases.²³ Another source of errors that cannot be adequately accounted for by qualitative methods deals with correlated (i.e., non-independent) variables. Correlation information can be applied in a quantitative context, given positive or negative correlation information, but a qualitative model cannot use this input information.²⁴ It is also significantly more difficult to empirically validate risk assessment models and methods that use ordinal scales.²⁵

2. Limitations in Determining Expected Value Risk with Ordinal Variables

The use of ordinal scales to measure variables (such as threat, vulnerability, and consequences) in risk assessments poses particular methodological problems if ordinal scale values for individual variables are combined (e.g., by multiplication) to form an overall risk

²⁰ D. V. Budescu, S. Broomell, and H.-H. Por, “Improving Communication of Uncertainty in the Reports of the Intergovernmental Panel on Climate Change,” *Psychological Science* 20 (2009): 299–308.

²¹ Cox, “What’s Wrong with Risk Matrices?”

²² Hubbard and Evans, “Problems with scoring methods and ordinal scales.”

²³ Ibid.

²⁴ Ibid.

²⁵ Ibid.

measure. As discussed above, it is not permissible to compute the mean (i.e., *average* or *expected value*) of a variable using a pure ordinal scale, and adding, subtracting, multiplying, or dividing pure ordinal scale variables is not meaningful. However, most individuals think in terms of ratio scales, where ratios of measurements are meaningful and a natural zero exists, and they have a tendency to improperly treat numerical ordinal scale measurements as being ratio-scaled.²⁶ For example, if Unlikely is assigned a value of 1 and Likely is assigned a value of 3, care must be taken to avoid improperly treating Likely events as actually being three times more likely to occur than Unlikely events.

If threshold ordinal scales are used, mean value computations are meaningful for the underlying ratio-scaled quantities, but the use of ordinal binning leads to a loss of information and interpretive ambiguity when the ordinal bins are used in expected value risk computations. The recommended solution, discussed in the main document accompanying this appendix, is to work directly with the ratio-scaled quantities in risk assessments and expected value calculations.

In the remainder of this appendix, an example that illustrates some of the problems that arise when variables measured using threshold ordinal scales are combined in expected value risk models is presented and discussed.

Suppose that the risk of an attack is modeled by the following simple function:

Risk = Threat x Consequences,

where, for a potential attack, Threat is the probability of the attack occurring and Consequences are a measure of the adverse impact given that the attack occurs (e.g., lives lost). This function is similar to other multiplicative models for risk, but is simplified for illustrative purposes. The Vulnerability variable has been removed and Risk is defined as the expected value of Consequences (i.e., as the expected loss). While this model is a simplified version of most risk assessments, the theoretical problems are analogous and can be illustrated with the simpler model. Suppose the variables are measured using the threshold ordinal scales in Table D-3.

If Threat and Consequences are assessed using these threshold ordinal scales and then combined via the above equation to compute Risk, all that can be deduced about a potential adverse event is that its Risk lies within a wide range. For example, if a particular event is assessed to have Medium Threat and Level 1 Consequences, the conclusion can only be that its expected Risk is somewhere between 200 and 5,000 lives lost. Additional information could potentially resolve some of this uncertainty and ambiguity, but such additional information is not captured by the ordinal scales for Threat and Consequences.

²⁶ Ibid.

Table D-3. Example Threshold Ordinal Scales

Ordinal Rating	Probability of Attack (Threat)	Ordinal Rating	Lives Lost (Consequences)
Low	0–0.2	Level 1	1000–10K
Medium	0.2–0.5	Level 2	10K–100K
High	0.5–1	Level 3	100K–1M
		Level 4	1M–10M

Table D-4. Ranges for Expected Value Risk Derived from Ordinal Scales

Ordinal Rating	Level 1	Level 2	Level 3	Level 4
Low	0–2K	0–20K	0–200K	0–2M
Medium	100–5K	2K–50K	20K–500K	200K–5M
High	500–10K	5K–100K	50K–1M	500K–10M

Using these scales leads to overlapping intervals for risk estimates, which in turns leads to an inability to determine the relative risk posed by different potential events. For instance, all Threat/Consequence level combinations overlap with the Low/Level 4 combination, so it is not possible to determine how the Low/Level 4 combination compares to any of the other possible combinations. Therefore, one cannot make any definitive statements about the relative risk of any combination compared to Low/Level 4 (e.g., it is not possible to determine whether an event assessed as Medium/Level 3 poses a higher or lower risk than an event assessed as Low/Level 4). The limited ability to determine relative magnitude of risks is not limited to the Low/Level 4 combination: the only comparisons that can be deduced regarding events assessed as Medium/Level 2 are that they pose more risk than Low/Level 1 events and less risk than Medium/level 4, High/Level 3, and High/Level 4 events.

Additional information about where Threat and Consequence variables lie within the ordinal bins (e.g., a precise point estimate of each variable on the underlying ratio scale or a probability distribution of possible values of the variable within the ordinal bin) would help resolve these overlapping issues and lead to better discrimination of the relative risks posed by potential adverse events, but such additional information is not captured by the ordinal scales. (If the underlying ratio scale values will be used, then the ordinal scales are unnecessary.) However, it should be noted that using a single point value for all variables in an ordinal bin (e.g., by assigning probability 0.35 to all Medium Threats) in an expected value risk calculation does not properly resolve the overlap issue; rather, it merely obscures it behind a false precision and may improperly change the assessments of key variables (e.g., an event assessed as having a Probability of Attack of 0.25 may be coded as Medium threat and as a result be improperly assigned a Probability of Attack of 0.35).

A common practice would be to combine the ordinal scales for Threat and Consequences into an overall risk matrix (instead of presenting the ranges for risk as in Table D-4). In this case, each pair of Threat and Consequence ordinal levels is assigned an overall ordinal risk level, e.g., from the set {Low, Medium, High, Very High}. Table D-5 gives an example of such a risk matrix that might be constructed from the ordinal scales in Table D-3. This risk matrix suffers from the same overlap and ambiguity issues discussed above, but it further compounds the problems by appearing to discriminate among risk ranges that actually overlap. For example, an event with Probability of Attack 0.001 that causes the loss of two million lives is considered High risk but an event with Probability of Attack 0.45 that causes the loss of 9,000 lives is considered Low risk, despite that fact that the expected number of lives lost is twice as great for the Low risk event). It is important to understand that the presence of such problems does not depend on the specific risk matrix chosen for this example. Although risk matrices can be of use to distinguish between very low and very high risks (e.g., the risk matrix in Table D-5 could be used to effectively distinguish between Low and Very High risks), in general it is not possible to construct risk matrices from threshold ordinal scales that properly discriminate risks in the middle ranges.²⁷ Moreover, these errors can be compounded when the underlying variables are correlated (i.e., not independent).²⁸

Table D-5. Example Risk Matrix Derived from Ordinal Scales

Ordinal Rating	Level 1	Level 2	Level 3	Level 4
Low	Low	Medium	Medium	High
Medium	Low	Medium	High	Very High
High	Medium	Medium	High	Very High

D. Proposals and Solutions

The use of ordinal scales to support risk assessments should be considered very carefully. When used correctly, they can be valuable for making quick, coarse comparisons and screening large numbers of events to separate very high risk from very low risk and to focus more attention and further analyses on higher risk events. While risk assessments based on ordinal scales are prevalent, their shortcomings, as outlined in this appendix, should not be overlooked. Consideration should be given to the potential sources of errors and biases associated with the use of ordinal scales (including ambiguity, range compression, and exacerbation of errors caused by cognitive biases and correlated variables). Variables measured with pure ordinal scales should not be multiplied to form a risk metric (the ordinal nature of the measurements implies that such

²⁷ Cox, “What’s Wrong with Risk Matrices?”

²⁸ Ibid.

multiplication is not meaningful), and variables measured with threshold ordinal scales should only be used in expected value risk computations with extreme caution.

Many approaches to risk formulate the risk metric as the product of probability and consequences, and compute consequences using ratio-scaled variables (e.g., dollars, lives). Probability estimates are intrinsically ratio-scaled and must obey the fundamental axioms and laws of probability theory. The resulting risk metric is often an expected value calculation (e.g., the expected loss in dollars or lives). The main document accompanying this appendix recommends this approach to estimating risk. Other possibilities, when this approach is used, include the use of Monte Carlo methods and event trees.²⁹ These methods are more rigorous analytically and can require more time, resources, and effort; in certain instances, this can be made more cost effective by combining the more rigorous quantitative methods with ordinal scaled screening methods.

Ordinal scales still maintain some utility for decision-makers, but only if the uncertainty of the risk calculation is clear and available to the decision-maker. For instance, when risk is calculated as an expected value from threshold ordinal scales, decision-makers should be aware of the potential range for that risk, so that they understand the strength of the analysis and can make defensible judgments. They should also be aware of the limitations on comparisons of risk that can be made, such as the inability to discriminate in the mid-ranges of risk comparisons. If the full extent of known uncertainty and methodological limitations are incorporated into a decision process, the resulting decision may appear initially to be less discriminating, but will, in fact, be more accurate and robust.

²⁹ J. Darrell Morgeson, Andrew J. Coe, and Victor A. Utgoff, *Review of Risk Assessment Methodologies for the Department of Homeland Security*, IDA Document D-3117 (Alexandria, VA: Institute for Defense Analyses, April 2005).

Appendix E

Illustrations

Figures

Figure 1. Types of Risk Informed Decisions that DHS Faces (in boxes) Arrayed Roughly According to the Decision-Making Horizon They Inform	8
Figure 2. Architecture for Managing a Scenario Space	30
Figure A-1. A Simple Layered Defense Model	A-4
Figure A-2. Illustration of Using Looking Up Tables for Probabilities of Success Generated by Expert Elicitation and Combining Them for the Overall Joint Probability of Success.....	A-8
Figure A-3. Notional Graphic Showing Layered Defense of Dams	A-9
Figure A-4. Integrated Process Control System	A-11
Figure A-5. Separated Process Control System	A-11
Figure A-6. Isolated Process Control System	A-12
Figure B-1. Notional Schematic of a Contagious Biological Agent Defense.....	B-7

Tables

Table 1. Sector-Specific Agencies and Assigned CIKR Sectors	11
Table 2. Asset and Systems-based Sectors	12
Table A-1. Layered Defense Configurations (LDCs) A–O	A-5
Table A-2. P(S A) Estimates for Each Attack Vector–LDC Combination.....	A-7
Table A-3. Cyber Defensive Configurations (CDCs) Based on Unique Sets of Cyber Defense Attributes.....	A-13
Table A-4. P(S A) Estimates for Each Attack Vector–CDC Combination	A-15
Table B-1. Weaponized Agent Parameters Summary Table	B-24
Table B-2. Target Population Parameters Summary Table	B-24
Table B-3. Target Physical Location Parameters Summary Table.....	B-25
Table B-4. Medical Intervention Parameters Summary Table	B-26
Table B-5. Policy-driven Intervention Parameters Summary Table.....	B-26
Table B-6. Steps for Contagious Biological Scenario Quantification	B-27
Table B-7. Initial Conditions Summary Table.....	B-28
Table B-8. Baseline Expected Value of Loss for Notional Scenario (Number of People Affected)	B-28

Table B-9. Baseline Expected Value of Loss for Notional Scenario with Medical Antibiotic Prophylaxis at Day 7 for 7 Days (Number of People Affected).....	B-29
Table B-10. Baseline Expected Value of Loss for Notional Scenario with Social Distancing at Day 7 (Number of People Affected).....	B-29
Table D-1. Common Examples of Measurement Scales	D-2
Table D-2. Summary Features of Measurement Scales.....	D-3
Table D-3. Example Threshold Ordinal Scales	D-9
Table D-4. Ranges for Expected Value Risk Derived from Ordinal Scales	D-9
Table D-5. Example Risk Matrix Derived from Ordinal Scales.....	D-10

Appendix F

References

- Atkinson, Leslie. 1988. "The Measurement Statistics Controversy: Factor Analysis and Subinterval Data." *Bulletin of the Psychonomic Society* 26: 261–264.
- Bauer, Timothy J., and Matthew G. Wolski. 2001. *Software User's Manual for the Chemical/Biological Agent Vapor, Liquid, and Solid Tracking (VLSTRACK) Computer Model, Version 3.1*. Dahlgren, VA: Naval Surface Warfare Center, Dahlgren Division.
- Coombs, C. H. 1950. "Mathematical Models in Psychological Scaling." *Journal of the American Statistical Association* 46: 480–489.
- Cooper, B. and D. P. Rice. 1976. "The Economic Cost of Illness Revisited." *Social Security Bulletin* 39: 21–36.
- Cox, Louis Anthony Jr. 2008. "What's Wrong with Risk Matrices?" *Risk Analysis* 28: 497–512.
- Cox, Louis Anthony Jr. 2005. "Some Limitations of Qualitative Risk Rating Systems." *Risk Analysis* 25: 651–662.
- Department of Energy. September. 16, 2008. *Risk Management Guide*. DOE G 413.3-7. Washington DC. <http://doe.test.doxcelerate.com/directives/archive-directives/413.3-EGuide-07/view>.
- Department of Homeland Security (DHS). 2006. *National Infrastructure Protection Plan*. Washington, DC: Department of Homeland Security.
- DHS. Risk Steering Committee. 2010. *DHS Risk Lexicon*, 2010 Edition. Washington DC: Department of Homeland Security.
- DHS. Office of Inspector General. June 2009. *Efforts to Identify Critical Infrastructure Assets and Systems*. DHS OIG-09-86. Washington, DC: Department of Homeland Security.
- Disraelly, D. S., T. J. Walsh, and C. A. Curling. 2011. "A New Methodology for Estimating Contagious Biological Agent Casualties as a Function of Time." *Mathematical and Computer Modelling* 54: 648–659.
- Franz, COL David R., Peter B. Jahrling, COL Arthur M. Friedlander, David J. McClain, COL David L. Hoover, COL W. Russell Bryne, MAJ Julie A. Pavlin, LT COL George W. Christopher, and COL Edward M. Eitzen, Jr. 1997. "Clinical Recognition and Management of Patients Exposed to Biological Warfare Agents." *Journal of the American Medical Association (JAMA)* 278 (5): 399–411.

- Gaito, John. May 1980. "Measurement Scales and Statistics: Resurgence of an Old Misconception." *Psychological Bulletin* 87 (3): 564–567.
- Henderson, Donald A., Thomas V. Inglesby, John G. Bartlett, Michael S. Ascher, Edward Eitzen, Peter B. Jahrling, Jerome Hauer, Marcelle Layton, Joseph McDade, Michael T. Osterholm, Tara O'Toole, Gerald Parker, Trish Perl, Philip K. Russell, MD, Kevin Tonat. 1999. "Smallpox as a Biological Weapon: Medical and Public Health Management." *JAMA* 281 (22): 2127–37.
- Hogarth, Robin M. 1975. "Cognitive Processes and the Assessment of Subjective Probability Distributions." *Journal of the American Statistical Association* 70: 271–289.
- Hubbard, Douglas W. 2009. *The Failure of Risk Management*. Hoboken, NJ: John Wiley & Sons.
- Hubbard D., and D. Evans. 2010. "Problems with scoring methods and ordinal scales in risk assessment." *IBM Journal of Research and Development* 54 (3): 1–10.
- Kampen, Jarl and Marc Swyngedouw. 2000. "The Ordinal Controversy Revisited." *Quality and Quantity* 34: 87–102.
- Landefeld, J. Steven and Eugene P. Seskin. 1982. "The Economic Value of Life: Linking Theory to Practice." *American Journal of Public Health* 72 (6): 555–566.
- Mandel, David R. 2005. "Are Risk Assessments of a Terrorist Attack Coherent?" *Journal of Experimental Psychology* 11: 277–288.
- Marcus-Roberts, Helen M., and Fred. S. Roberts. 1987. "Meaningless Statistics." *Journal of Educational Statistics* 12 (4): 383–394.
- Morgeson, J. Darrell and Peter Brooks. 2011. *DHS Risk Assessment*. IDA Draft Paper. Alexandria, VA: Institute for Defense Analyses.
- Morgeson, Darrell J., Andrew J. Coe, and Victor A. Utgoff. April 2005. *Review of Risk Assessment Methodologies for the Department of Homeland Security*. IDA Document D-3117. Alexandria, VA: Institute for Defense Analyses.
- National Academy of Sciences. National Research Council. 2010. *Review of the Department of Homeland Security's Approach to Risk Analysis*. Washington, DC: National Academies Press.
- North Atlantic Treaty Organization (NATO). 2010. *Allied Medical Publication 8(C)—North Atlantic Treaty Organization (NATO) Planning Guide for the Estimation of Chemical, Biological, Radiological, and Nuclear (CBRN) Casualties*, Ratification Draft (AMedP-8(C)). DRAFT.
- Siegel, S. 1957. "Nonparametric Statistics." *The American Statistician* 11: 13–19.
- Singpurwalla, Nozer D. 2006. *Reliability and Risk: A Bayesian Perspective*. Chichester, England: John Wiley & Sons, Ltd.
- Stevens, S. S. 1951. "Mathematics, measurement and psychophysics." In S.S. Stevens (Ed.), *Handbook of Experimental Psychology*. New York: John Wiley & Sons, Inc.
- Stevens, S. S. 1946. "On the theory of scales of measurement." *Science* 103: 677–680.

- Stoneburner, Gary, Alice Goguen, and Alexis Feringa. July 2002. *Risk Management Guide for Information Technology Systems*. NIST 800-30. Gaithersburg, MD: National Institute of Standards and Technology.
- Titan Corporation and Defense Threat Reduction Information Analysis Center. 2004. *Hazard Prediction Assessment Capability (HPAC) Getting Started*. Virginia: Defense Threat Reduction Agency.
- Velleman, Paul and Leland Wilkinson. 1993. “Nominal, Ordinal, Interval, and Ratio Typologies Are Misleading.” *The American Statistician* 47 (1): 65–72.
- Wason, P.C. 1960. “On the failure to eliminate hypotheses in a conceptual task.” *The Quarterly Journal of Experimental Psychology* 12: 129–140.
- Wilson, Thomas P. 1971. “Critique of Ordinal Variables.” *Social Forces* 49: 432–44.
- Wu, Lien-The. 1926. “A Treatise on Pneumonic Plague,” C.H.474. Geneva: League of Nations Health Organization.

Appendix G

Abbreviations

P(A)	Probability that a Scenario Occurs
C	Consequences
C-AV	Cyber Attack Vector
CB Strike	Chemical/biological strike
CBRN	Chemical, Biological, Radiological, and Nuclear
CDC	Cyber Defensive Configuration
CIKR	Critical Infrastructure and Key Resources
CNCI	Comprehensive National Cybersecurity Initiative
COTS	Commercial Off-the-Shelf
DHS	Department of Homeland Security
DNI	Director of National Intelligence
DOD	Department of Defense
DTM	Directive-Type Memorandum
EVC	Expected Value of the Consequences
FBI	Federal Bureau of Investigation
GIDEP	Government-Industry Data Exchange Program
HPAC	Hazard Prediction and Assessment Capability
HRIP	Human Response Injury Profile
HSPD	Homeland Security Presidential Directive
IATAC	Information Assurance Technology Analysis Center
ICT	Information and Communications Technology
IDA	Institute for Defense Analyses
IED	Improvised Explosive Device
IT	Information Technology
JAMA	Journal of the American Medical Association
LDC	Layered Defense Configuration
NATO	North Atlantic Treaty Organization
NFC	Near Field Communications
NIPP	National Infrastructure Protection Plan
NISAC	National Infrastructure Simulation and Analysis Center
NRC	National Research Council
OCI	Office of Counterintelligence
OEM	Original Equipment Manufacturer
OTS	Off-the-Shelf
P	Probability
PCS	Process Control System
PAR	Population at risk
PRC	People's Republic of China
P(S A)	Probability that an Attack is Successful
R	Risk

R&D	Research and Development
RFID	Radio Frequency Identification
ROI	Return on Investment
SARS	Severe Acute Respiratory Syndrome
SCADA	Supervisory Control and Data Acquisition
SEIRP	Susceptible, Exposed and Infected, Infectious, Removed, and Prophylaxis efficacious
SOAR	State of the Art Report
SSP	Sector-Specific Plan
T	Threat
TCP/IP	Transmission Control Protocol/Internet Protocol
U.S.	United States
USB	Universal Serial Bus
USG	United States Government
V	Vulnerability
VBIED	Vehicle-Bourne Improvised Explosive Device
VLSTRACK	Vapor, Liquid, Solid Tracking

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YY) December 2011		2. REPORT TYPE Final		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE Doctrinal Guidelines for Quantitative Vulnerability Assessments of Infrastructure-Related Risks, Volume I				5a. CONTRACT NO. DASW01-04-C-0003	
				5b. GRANT NO.	
				5c. PROGRAM ELEMENT NO(S).	
6. AUTHOR(S) J. Darrell Morgeson, Project Leader, Peter S. Brooks, Deena S. Disraelly, Jeremy L. Erb, Michael L. Neiman, Whitney C. Picard				5d. PROJECT NO. ER-6-2474.02	
				5f. WORK UNIT NO.	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Institute for Defense Analyses 4850 Mark Center Drive Alexandria, VA 22311-1882				8. PERFORMING ORGANIZATION REPORT NO. IDA Document D-4477, Volume I	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Department of Homeland Security 245 Murray Lane, SW Washington, DC 20528				10. SPONSOR'S / MONITOR'S ACRONYM(S) DHS	
				11. SPONSOR'S / MONITOR'S REPORT NO(S).	
12. DISTRIBUTION / AVAILABILITY STATEMENT					
13. SUPPLEMENTARY NOTES Approved for public release; distribution is unlimited.					
14. ABSTRACT The objective of this document is to provide doctrinal guidelines for operationalizing a framework for quantifying risk, with a specific focus on quantitatively estimating the vulnerability of assets and systems comprising the nation's critical infrastructure. IDA focused on vulnerability for three reasons. First, its definition and how it is applied to critical infrastructure is far less understood than the concepts of threat and consequence. Second, a sound approach for quantifying vulnerability will improve the methodologies for quantifying risk for critical infrastructure. Third, clearly defining vulnerability is key to developing commensurate risk metrics across the 18 critical infrastructure and key resources (CIKR) sectors. When systems vulnerability and asset vulnerability protected by layered defenses are compared side-by-side, the overall recommendation is to define vulnerability as the expected value of loss given a scenario occurrence in both cases. This requires that vulnerability for layered defenses be re-interpreted as the product of the joint probability of successfully penetrating all relevant defensive layers, and consequences. IDA sought to define a set of concepts and computational methods for quantifying vulnerability in a way that the resulting risk calculations produce commensurable risk metrics regardless of whether the risks are related to systems or isolated assets, or due to natural hazards or adversarial threats.					
15. SUBJECT TERMS Risk, Threat, Vulnerability, Consequences, CIKR					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NO. OF PAGES 114	19a. NAME OF RESPONSIBLE PERSON
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include Area Code)

